arXiv:1612.05222v2 [cs.DS] 19 Dec 2016

Multi-Agent and Multivariate Submodular Optimization

Richard Santiago[1] and F. Bruce Shepherd[2]

[1] McGill University, Canada, richard.santiagotorres@mail.mcgill.ca
[2] McGill University, Canada, bruce.shepherd@mcgill.ca

Abstract. Recent years have seen many algorithmic advances in the area of submodular optimization: (SO) $\min / \max\ f(S) : S \in \mathcal{F}$, where $\mathcal{F}$ is a given family of feasible sets over a ground set V and $f : 2^V \to \mathbb{R}$ is submodular. This progress has been coupled with a wealth of new applications for these models. Our focus is on a more general class of *multi-agent submodular optimization* (MASO) which was introduced by Goel et al. in the minimization setting: $\min \sum_i f_i(S_i) : S_1 + S_2 + \cdots + S_k \in \mathcal{F}$. Here we use $+$ to denote disjoint union and hence this model is attractive where resources are being allocated across k agents, each with its own submodular cost function $f_i()$. In this paper we explore the extent to which the approximability of the multi-agent problems are linked to their single-agent *primitives*.

We present a generic reduction that transforms a multi-agent (and more generally multivariate) problem into a single-agent one, showing that several properties of the objective function and family of feasible sets are preserved. For maximization, this allows one to leverage algorithmic results from the single-agent submodular setting for the multi-agent case. We are not aware of work for general families in the maximization setting and so these results substantially expand the family of tractable models. For instance, we see that multi-agent (and multivariate) maximization subject to a p-matroid constraint can be reduced to a single-agent problem over the same type of constraint (i.e., p matroid intersections). Allowing the (much) more general class of multivariate submodular objective functions $f(S_1, S_2, \ldots, S_k)$ (instead of decomposable functions $f = \sum_i f_i$) gives new modelling capabilities. We describe a family of multivariate submodular objectives based on quadratic functions and show how these may be used to penalize competition between agents.

For monotone multi-agent minimization we give three types of reductions. The first two are black-box reductions which show that (MASO) has an $O(\alpha \cdot \min\{k, \log^2(n)\})$-approximation when (SO) over $\mathcal{F}$ admits an $O(\alpha)$-approximation over the natural formulation. The third type of reduction is for the special case when every element in the blocking clutter of $\mathcal{F}$ has size at most β. Here we show that β-approximations and $\beta \ln(n)$-approximations are always available for the single-agent and multi-agent problems respectively. Moreover, virtually all known approximations for single-agent monotone submodular minimization are derived in this way.

1 Introduction

For a family of feasible sets $S \in \mathcal{F}$ on a finite ground set V we consider the following broad class of submodular optimization (SO) problems:

$$\text{SO}(\mathcal{F}) \qquad \text{Min} \ / \ \text{Max} \ \ f(S) : S \in \mathcal{F} \tag{1}$$

where f is a submodular set function on V. We call these two problems the "primitives" associated with $\mathcal{F}$. There has been an impressive recent stream of activity around these

problems for various set families $\mathcal{F}$. We explore the connections between these primitive problems and their multi-agent incarnations. In the *multi-agent (MA)* version, we have k agents each of which has an associated submodular set function f_i, $i \in [k]$. As before, we are looking for sets $S \in \mathcal{F}$, however, we now have a 2-phase task: the elements of S must also be partitioned amongst the agents. Hence we have set variables S_i and seek to optimize $\sum_i^k f_i(S_i)$. This leads to the multi-agent submodular optimization (MASO) versions:

$$\text{MASO}(\mathcal{F}) \qquad \text{Min} \;\; / \;\; \text{Max} \;\; \textstyle\sum_i f_i(S_i) : S_1 + S_2 + \cdots + S_k \in \mathcal{F}. \tag{2}$$

The special case when $\mathcal{F} = \{V\}$ was previously examined but the general form was only introduced and developed in Goel et al. [12] for the minimization setting. A natural first question is whether any multi-agent problem could be reduced (or encoded) to a single-agent one. Goel et al. show this is not the case. The following simple submodular-cost problem has a MA cost function which cannot be realized by a single submodular cost function over the same ground set V. There are three tasks $\{A, B, C\}$ and two agents $\{1, 2\}$. Each agent can do one of the tasks at cost 1 each. In addition agent 1 can do tasks $\{A, B\}$ at cost of only 1, and $\{A, B, C\}$ at cost of only 2. Agent 2 can do $\{A, C\}$ at cost of only 1, and $\{A, B, C\}$ at cost of only 2. We can associate a cost $c(S)$ for any subset $S \subseteq \{A, B, C\}$ which corresponds to the best allocation of tasks in S to the agents. This function $c(S)$ is not, however, submodular even though each agent's cost function was submodular. To see this, note that the marginal cost of adding B to $\{A\}$ is 0 which is smaller than the marginal cost of adding it to $\{A, C\}$.

Perhaps even more emphatically, they prove that when $\mathcal{F}$ consists of vertex covers in a graph, the *single-agent (SA)* version (i.e., (1)) has a 2-approximation while the MA version has an inapproximability lower bound of $\Omega(log(n))$. One of our main objectives is to explain the extent to which approximability for multi-agent problems is intrinsically connected to their single-agent versions.

In some cases our methods get traction for a more general class of problems that we call the *capacitated multi-agent submodular optimization (CMASO) problem*

$$\text{CMASO}(\mathcal{F}) \qquad \begin{array}{rl} \max / \min & \sum_{i=1}^k f_i(S_i) \\ \text{s.t.} & S_1 + \cdots + S_k \in \mathcal{F} \\ & S_i \in \mathcal{F}_i \,, \forall i \in [k] \end{array} \tag{3}$$

where we are supplied with the subfamilies $\mathcal{F}_i$. Many existing applications fit into this framework and some of these can be enriched through the added flexibility of the capacitated model. For instance, one may include set bounds on the variables: $L_i \subseteq S_i \subseteq U_i$ for each i, or simple cardinality constraints: $|S_i| \leq b_i$ for each i. Applications which fit into this framework are discussed further in Section 1.2.

1.1 Multivariate Submodular Optimization

Multi-agent optimization can be seen as a special case of a much more general class of multivariate submodular optimization (MVSO) problems. Namely functions of several variables which satisfy some type of submodularity properties. In this paper, we call a multivariate function $f : (2^V)^k \to \mathbb{R}$ *multi-submodular* (or sometimes multivariate submodular) if for all pairs of tuples $(S_1, ..., S_k), (T_1, ..., T_k) \in (2^V)^k$ we have

$$f(S_1, ..., S_k) + f(T_1, ..., T_k) \geq f(S_1 \cup T_1, S_2 \cup T_2, .., S_k \cup T_k) + f(S_1 \cap T_1, ..., S_k \cap T_k).$$

We mention that other extensions have been considered in the literature, including [46,43,20,23]. Notice that in the special case of $k = 1$ a multi-submodular function corresponds to a univariate submodular function. Hence, the above can be thought as a natural extension of submodularity to multivariate functions. Since the objective functions used in CMASO are multi-submodular, this leads to a more general class of problems:

$$\text{CMVSO}(\mathcal{F}) \qquad \begin{array}{c} \max / \min \ f(S_1, S_2, \ldots, S_k) \\ \text{s.t. } S_1 + \cdots + S_k \in \mathcal{F} \\ S_i \in \mathcal{F}_i \,, \forall i \in [k]. \end{array} \tag{4}$$

While several of our multi-agent results extend to the multivariate setting, the class of multi-submodular functions can encode much more than the "decomposable" functions arising in multi-agent objectives. We give complexity-theoretic evidence of this in Section 5 where we establish (see Corollary 11 and Lemma 6) a large gap between the following two problems.

$$(\mathbf{MV-Min}) \quad \begin{array}{c} \min \quad g(S_1, \ldots, S_k) \\ \text{s.t. } S_1 + \cdots + S_k = V \\ S_i \subseteq V_i \\ |S_i| \leq 1 \end{array} \qquad (\mathbf{MA-Min}) \quad \begin{array}{c} \min \quad \sum_{i \in [k]} g_i(S_i) \\ \text{s.t. } S_1 + \cdots + S_k = V \\ S_i \subseteq V_i \\ |S_i| \leq 1 \end{array}$$

Proposition 1. *The problem (MV-Min) is $\Omega(n)$-inapproximable under the oracle model (even if g is monotone and nonnegative) whereas the problem (MA-Min) has an exact polytime algorithm (for arbitrary submodular functions g_i).*

The remainder of the introduction is structured as follows. In Section 1.3, we describe our generic reductions which turn MA (and in some cases, the more general multivariate) optimization problems into a SA primitive. This reduction together with recent algorithmic advances in (SA) submodular optimization imply that many interesting instances of CMASO become tractable. Prior work in both the single and multi-agent settings is summarized in Section 1.4. We believe that the CMASO and CMVSO frameworks have potential for future applications and so we start by presenting some existing examples which fit into, and can be further extended, in these models.

1.2 Some Applications of Capacitated Multi-Agent and Multivariate Optimization

In this section we present several problems in the literature which are special cases of Problem (3). We also indicate cases where the extra generality of CMASO gives modelling advantages. We postpone the algorithmic and complexity results until Section 1.4. We start by discussing the maximization setting.

Example 1 (The Submodular Welfare Problem). The most basic problem in the maximization setting arises when we simply take the feasible space $\mathcal{F} = \{V\}$. This describes a well-known model (introduced in [31]) for allocating goods from a ground set to agents, each of which has a submodular valuation (utility) function over baskets of goods. This is formulated as (3) by considering nonnegative monotone functions f_i, $\mathcal{F} = \{V\}$, and $\mathcal{F}_i = 2^V$ for all i. The computational problem is then

$$\begin{array}{c} \max \quad \sum_{i=1}^{k} f_i(S_i) \\ \text{s.t. } S_1 + \cdots + S_k = V. \end{array}$$

Example 2 (Sensor Placement with Multivariate Objectives). The problem of placing sensors and information gathering has been popular in the submodularity literature [24,27,25]. We are given a set of sensors V and a set of possible locations $\{1, 2, \ldots, k\}$ where the sensors can be placed. There is also a budget constraint restricting the total number of sensors that can be deployed. The goal is to place sensors at some of the locations so as to maximize the "informativeness" gathered. We argue below that this application is well suited to a multivariate submodular objective function $f(S_1, ..., S_k)$ which measures the "informativeness" of placing sensors S_i at location i. A natural mathematical formulation for this is given by

$$\begin{array}{ll} \max & f(S_1, ..., S_k) \\ \text{s.t.} & S_1 + \cdots + S_k \in \mathcal{F} \\ & S_i \in \mathcal{F}_i, \end{array}$$

where $\mathcal{F} := \{S \subseteq V : |S| \leq b\}$ imposes the budget constraint and $\mathcal{F}_i$ gives additional modelling flexibility for the types (or number) of sensors that can be placed at location i. For instance, we could impose a cardinality constraint $|S_i| \leq b_i$ by defining $\mathcal{F}_i = \{S \subseteq V : |S| \leq b_i\}$. We could also impose that only sensors $V_i \subseteq V$ can be placed at location i by taking $\mathcal{F}_i = \{S \subseteq V : S \subseteq V_i\}$. We can also combine these constraints by taking $\mathcal{F}_i = \{S \subseteq V_i : |S| \leq b_i\}$. Notice that in these cases both $\mathcal{F}$ and the $\mathcal{F}_i$ are matroids and hence our algorithms will apply. One may form a multivariate objective by defining $f(S_1, S_2, \ldots, S_n) = \sum_i f_i(S_i) - R(S_1, S_2, \ldots, S_n)$ where the f_i's measure the benefit of placing sensors S_i at location i, and $R()$ is a redundancy function. If the f_i's are submodular and $R()$ is multivariate supermodular, then f is multivariate submodular. In this setting, it is natural to take the f_i's to be coverage functions (i.e. $f_i(S_i)$ measures the coverage of placing sensors S_i at location i). We next propose a family of "redundancy" functions which are supermodular.

Supermodular penalty measures via Quadratic functions. We denote $\mathbf{S} := (S_1, S_2, \ldots, S_n)$ and define $z_{\mathbf{S}} := (|S_1|, |S_2|, \ldots, |S_n|)$. One can show (see Lemma 14 in Appendix C) that if A is a matrix satisfying $a_{ij} + a_{ji} \geq 0$, then $R(\mathbf{S}) := z_{\mathbf{S}}^T A z_{\mathbf{S}}$ is multivariate supermodular. For instance, in the example, one could take redundancy coefficients a_{ij} as $\Theta(\frac{1}{d(i,j)^2})$ where $d(i, j)$ denotes the distance between locations i and j. This can be further extended so that different sensor types contribute different weights to the vector $z_{\mathbf{S}}$, e.g., define $z_{\mathbf{S}}(i) = \sum_{j \in S_i} w(j)$ for an associated sensor weight vector w. We believe that these quadratic functions (based on matrices satisfying $a_{ij} + a_{ji} \geq 0$) yield a useful class of multivariate submodular functions for modelling "competition" between agents in many domains other than sensor problems.

Example 3 (Recommendation Systems and Matroid Constraints). This has been a widely deployed class of problems that usually entails the targetting of product ads to a mass of (largely unknown) buyers or "channels". In [5] a "meta" problem is considered where (known) prospective buyers are recommended to interested sellers. This type of recommendation system incurs additional constraints such as (i) bounds on the size of the buyer list provided to each seller (possibly constrained by a seller's advertising budget) and (ii) bounds on how often a buyer appears on a list (so as to not bombard them). They model these constraints as a "b-matching" problem in a bipartite buyer-seller graph G_B. They also consider a more exotic model which incorporates "conflict-aware" constraints on the buyer list for each seller (e.g., no more than two buyers from the same household). Conflicts are modelled by extra edges amongst the buyer nodes: the subgraph of buyers recommended to a seller has an upper bound on the number of allowed conflict edges. Heuristics for this (linear-objective)

model [5] are successfully developed on Ebay data, even though the computational problem was shown to be NP-hard. In fact, subsequent work [4] shows that conflict-aware b-matching is unlikely to have a polytime algorithm which approximates the optimum to within factors of $O(n^{1-\epsilon})$. We now propose an alternative model which admits an $O(1)$-approximation.

For sellers with diminishing returns, a natural alternative to a linear objective function is a decomposable multivariate submodular objective $f(S_1, S_2, \ldots, S_k) = \sum_i f_i(S_i)$ where S_i denotes the list of buyers recommended to seller i. (Multivariate supermodular penalty functions, discussed in the previous example, may also be used to tap into competition between sellers.) The model (3) can now be used to reformulate the conflict constraints so as to avoid the aforementioned polynomial inapproximability. We consider the problem

$$\begin{array}{ll} \max & \sum_{i=1}^{k} f_i(S_i) \\ \text{s.t.} & S_1 \cup S_2 \ldots \cup S_k \in \mathcal{F} \\ & S_i \in \mathcal{F}_i, \ \ \forall i \in [k] \end{array}$$

where $\mathcal{F}$ is a partition matroid on the buyers enforcing the upper bound on the number of times a buyer is recommended. We define $\mathcal{F}_i$ to enforce constraints for the seller. First, we may restrict the family to the feasible buyers $V_i \subseteq V$ for each seller i. Second we may partition V_i into "households" V_{ij} and require that seller i's list include at most 1 element from each V_{ij}. Thus $\mathcal{F}_i$ is a partition or laminar matroid. Our results will show that this version has a polytime $O(1)$-approximation (in the oracle model).

Example 4 (Public Supply Chain). Consider the problem of transporting resources from the hinterland to external markets served by a single port. As prices fluctuate the underlying value of resources changes and this motivates the expansion of public or private networks. These problems have a large number of stakeholders such as port authority, multiple train and trucking carriers, labor unions and overseeing government transportation bodies. It is not uncommon for one good (say coal) to be transported by several distinct organizations and even more than one firm involved in the same mode of transport (say rail). The design and expansion of such networks therefore represent examples of strong multi-agent behaviour. Moreover, feasible sets based on matroid constraints lend themselves to very natural transportation structures such as arborescences (which are a special case of intersection of two matroids).

We now discuss Problem (3) in the minimization setting.

Example 5 (Minimum Submodular Cost Allocation). Similar to the maximization setting, the most basic problem in the minimization setting arises when we simply take the feasible space $\mathcal{F} = \{V\}$, and the agents are allowed to take any subset of V (i.e. $\mathcal{F}_i = 2^V$ for all i). This problem,

$$\begin{array}{ll} \min & \sum_{i=1}^{k} f_i(S_i) \\ \text{s.t.} & S_1 + \cdots + S_k = V, \end{array}$$

has been widely considered in the literature for both monotone and nonmonotone functions [42,7,3], and is referred to as the Minimum Submodular Cost Allocation (MSCA) problem[3] (introduced in [17,42] and further developed in [3]).

The framework from (3) allows us to incorporate additional constraints into the MSCA problem by defining the families $\mathcal{F}_i$ appropriately. The most natural are to impose cardinality

[3] Sometimes referred to as submodular procurement auctions.

constraints on the number of elements that an agent can take, or to only allow agent i to take a set S_i of elements satisfying some bounds $L_i \subseteq S_i \subseteq U_i$. For instance, in the Submodular Facility Location problem, one can model that a potential facility must serve some customers, and may only serve customers within a pre-specified region given by U_i. These are discussed in detail on Section 5.

Example 6 (Multi-agent Minimization). Goel et al [12] consider the case of Problem (3) for monotone, submodular functions in which $\mathcal{F}$ is a nontrivial collection of subsets of V (i.e. $\mathcal{F} \subset 2^V$) and there is no restriction on the $\mathcal{F}_i$ (i.e. $\mathcal{F}_i = 2^V$ for all i). In particular, given a graph G they consider the families of vertex covers, spanning trees, perfect matchings, and shortest st paths.

Example 7 (Rings). It is known [40] that arbitrary submodular functions can be minimized efficiently on a ring family, i.e. a family of sets closed under unions and intersections (see [38] for faster implementations, and work on modified ring families [14]). A natural extension is to minimization of multivariate submodular functions over a *multivariate ring family,* where by the latter we mean a family of tuples closed under (component wise) unions and intersections. We provide the formal definitions in Section 2.4 and show that this more general problem can still be solved efficiently by applying the reduction from Section 2.

1.3 Our Contributions

We discuss a generic reduction that transforms the multivariate problem (4) (and hence (3)) into a single-agent problem. The reduction uses the simple idea of viewing assignments of elements v to agents i in a *multi-agent bipartite graph* (this was first proposed by Nisan et al in [32] and also used in Vondrak's work [44]; in these cases $\mathcal{F} = \{V\}$). The associated bipartite graph is $G([k], V) = ([k] + V, E)$ where there is an edge iv for each $i \in [k]$ and each $v \in V$. Any valid assignment of some elements in V to agents corresponds to a b-matching Z in G, where $b[i] = \infty$ for each $i \in [k]$ and $b(v) = 1$ for each $v \in V$.

While the above auxiliary graph is fine when $\mathcal{F} = \{V\}$ (as in [32,44]), in general multi-agent optimization problems we may use a subset $F \subset V$ from a general family $\mathcal{F}$. Hence we are seeking b-matchings Z for which $cov(Z) \in \mathcal{F}$. Here we define $cov(Z)$ to be the set of nodes $v \in V$ which are saturated by Z. There is no apriori reason that these *$\mathcal{F}$-constrained b-matchings* should have any nice structure, even for simple $\mathcal{F}$. Our first result (Section 2) is that indeed this is the case if $\mathcal{F}$ is a matroid.

Theorem 1. *If $\mathcal{F}$ induces a matroid over V, then so does the family of $\mathcal{F}$-constrained b-matchings over the edges of the multi-agent bipartite graph G.*

Moreover, nice structure is also maintained if we start with a family of matroid bases or an intersection of p matroids. This, together with the recent advances in SA submodular optimization (Section 1.4), yields the following algorithmic consequences.

Theorem 2. *Consider the capacitated multivariate* **maximization** *problem*

$$\begin{aligned} \max \;\; & f(S_1, S_2, \ldots, S_k) \\ s.t. \;\; & S_1 + \cdots + S_k \in \mathcal{F} \\ & S_i \in \mathcal{F}_i \,, \forall i \in [k] \end{aligned}$$

where f is a nonnegative multivariate submodular function and the families $\mathcal{F}_i$ are matroids. Then we have the following results.

1. *If $\mathcal{F}$ is a matroid, then there is a $(2+\epsilon)$-approximation for monotone functions and a $(2+O(1))$-approximation for nonmonotone functions. Moreover, in the special case where $\mathcal{F}_i = 2^V$ for all i, there is a $(1-\frac{1}{e})$-approximation for monotone functions and a $\frac{1}{e}$-approximation for nonmonotone functions.*
2. *If $\mathcal{F}$ is a p-matroid intersection, then there is a $(p+1+\epsilon)$ and a $(p+1+O(1))$-approximation for monotone and nonmonotone functions respectively. Moreover, in the case where $\mathcal{F}_i = 2^V$ for all i, there is a $(p+\epsilon)$ for monotone functions.*
3. *If $\mathcal{F}$ is the set of bases of a matroid and $\mathcal{F}_i = 2^V$ for all i, then there is a $(1-\frac{1}{e})$-approximation for monotone functions and a $\frac{1}{2}(1-\frac{1}{\nu})$-approximation for nonmonotone functions where ν denotes the fractional base packing number.*

The nice properties of the reduction also allow us to leverage results for single-agent robust submodular maximization to the multivariate setting. The robustness considered in this case is with respect to the adversarial removal of up to τ elements. We point out that while this paper does not focus on robust formulations, we believe that this result may be of independent interest and could lead to interesting applications and future work. We prove the following.

Theorem 3. *Consider the multivariate* **robust maximization** *problem*

$$\begin{array}{cc} \max & \min \\ S_1+\cdots+S_k \in \mathcal{F} & A_i \subseteq S_i \\ S_i \in \mathcal{F}_i & \sum_i |A_i| \leq \tau \end{array} \quad f(S_1 - A_1, S_2 - A_2, \ldots, S_k - A_k),$$

where f is a nonnegative monotone multivariate submodular function and the families $\mathcal{F}_i$ are matroids. Then we have the following.

1. *There is a $[(\tau+1)(p+1+\epsilon)]$-approximation for any fixed $\epsilon > 0$ in the case where $\mathcal{F}$ is a p-matroid intersection.*
2. *In the special case where $\mathcal{F}_i = 2^V$ for all i, there is a $(\frac{\tau+1}{1-1/e})$-approximation in the case where $\mathcal{F}$ is a matroid, and a $[(\tau+1)(p+\epsilon)]$-approximation for any fixed $\epsilon > 0$ in the case where $\mathcal{F}$ is a p-matroid intersection.*

On the minimization side, there exist polynomial inapproximability results even in the single-agent setting subject to a matroid constraint (e.g. cardinality constraint [41] or spanning trees [12]). Hence the "good behaviour" of the reduction (Theorem 1) is not directly applicable. However, we describe two different connections (in terms of n and in terms of k) between MA and SA optimization. These are based on the natural "blocking" LP relaxation; see Appendices A and B.

Theorem 4. *Suppose there is a (polytime) $\alpha(n)$-approximation for the single-agent nonnegative monotone submodular minimization problem based on rounding the natural LP for a family $\mathcal{F}$. Then there is a (polytime) $O(\alpha(n) \cdot \min\{k, \log^2(n)\})$-approximation for the multi-agent nonnegative monotone submodular minimization problem over $\mathcal{F}$. Moreover, the $O(k\alpha(n))$-approximation extends to the more general multivariate setting.*

In Section 4, we revisit the state for constrained minimization primitives and spotlight some new cases of interest which have good approximability properties. Given a family $\mathcal{F}$ the *blocker* $\mathcal{B}(\mathcal{F})$ of $\mathcal{F}$ consists of the minimal sets B such that $B \cap F \neq \emptyset$ for each $F \in \mathcal{F}$. We say that $\mathcal{B}(\mathcal{F})$ is β-*bounded* if $|B| \leq \beta$ for all $B \in \mathcal{B}(\mathcal{F})$. Our results can be summarized as follows.

Theorem 5. *Let $\mathcal{F}$ be a family with a β-bounded blocker. Then there are β-approximation and $\beta \ln(n)$-approximation algorithms for the associated single-agent and multi-agent non-negative monotone problems respectively. These yield polytime algorithms if $P^*(\mathcal{F})$ has a polytime separation oracle.*

Corollary 1. *If $\mathcal{B}(\mathcal{F})$ is β-bounded and there is a polytime separation oracle for the blocking formulation of $\mathcal{F}$, then there is a polytime β-approximation for the above minimization problem.*

We mention that this result implies the 2-approximation for vertex cover, and k-approximation for k-uniform hitting set, which are the only "good" approximations (known to us) for the minimization version of (1) for nontrivial families $\mathcal{F}$ [12,19].

1.4 Related Work

Single Agent Optimization. The high level view of the tractability status for unconstrained (i.e., $\mathcal{F} = 2^V$) submodular optimization is that both maximization and minimization are generally behave well. Minimizing a submodular set function is a classical combinatorial optimization problem which can be solved in polytime [15,40,18]. Unconstrained maximization on the other hand is known to be inapproximable for general submodular set functions but admits a polytime constant-factor approximation algorithm when f is nonnegative [1,9].

In the constrained maximization setting, the classical work [36,37,11] already established an optimal $(1 - \frac{1}{e})$-approximation factor for maximizing a nonnegative monotone submodular function subject to a cardinality constraint, and a $(k + 1)$-approximation for maximizing a nonnegative monotone submodular function subject to k matroid constraints. This approximation is almost tight in the sense that there is an (almost matching) factor $\Omega(k/\log(k))$ inapproximability result. For nonnegative monotone functions [44,2] give an optimal $(1 - \frac{1}{e})$-approximation based on multilinear extensions when $\mathcal{F}$ is a matroid; [28] provides a $(1 - \frac{1}{e} - \epsilon)$-approximation when $\mathcal{F}$ is given by a constant number of knapsack constraints, and [30] gives a local-search algorithm that achieves a $(k + \epsilon)$-approximation (for any fixed $\epsilon > 0$) when $\mathcal{F}$ is a k-matroid intersection. For nonnegative nonmonotone functions, a $\frac{1}{e}$-approximation is known [10] for maximization under a matroid constraint, in [29] a $(k + O(1))$-approximation is given for k matroid constraints with k fixed. A simple "multi-greedy" algorithm [16] matches the approximation of Lee et al. but is polytime for any k. Finally, Vondrak [45] gives a $\frac{1}{2}(1-\frac{1}{\nu})$-approximation under a matroid base constraint where ν denotes the fractional base packing number.

Recently, Orlin et al [39] considered a robust formulation for constrained monotone submodular maximization, previously introduced by Krause et al [26]. The robustness in this model is with respect to the adversarial removal of up to τ elements. The problem can be stated as $\max_{S\in\mathcal{F}} \min_{A\subseteq S, |A|\leq\tau} f(S - A)$, where $\mathcal{F}$ is an independence system. Assuming an α-approximation is available for the special case $\tau = 0$ (i.e. the non-robust version), [39] gives an $(\frac{\alpha}{\tau+1})$-approximation for the robust problem.

For constrained minimization, the news are worse [12,41,19]. For instance, if $\mathcal{F}$ consists of spanning trees (bases of a graphic matroid) Goel et al [12] show a lower bound of $\Omega(n)$, while in the case where $\mathcal{F}$ corresponds to the cardinality constraint $\{S : |S| \geq k\}$ Svitkina and Fleischer [41] show a lower bound of $\tilde{\Omega}(\sqrt{n})$. There are a few exceptions. The problem can be solved exactly when $\mathcal{F}$ is a ring family ([40]), triple family ([15]), or parity family ([13]). In the context of NP-Hard problems, there are almost no cases where good (say $O(1)$

or $O(\log(n))$) approximations exist. We have that the submodular vertex cover admits a 2-approximation ([12,19]), and the k-uniform hitting set has $O(k)$-approximation.

Multi-agent Problems. In the maximization setting the main multi-agent problem studied is the Submodular Welfare Maximization ($\mathcal{F} = \{V\}$) for which the initial $\frac{1}{2}$-approximation [32] was improved to $(1 - \frac{1}{e})$ by Vondrak [44] who introduced the continuous greedy algorithm. This approximation is in fact optimal [21,34]. We are not aware of maximization work for Problem (3) for a nontrivial family $\mathcal{F}$.

For the multi-agent minimization setting, MSCA is the most studied application of Problem (3) ($\mathcal{F} = \{V\}$). If the functions f_i are nonnegative and nonmonotone, [7] recently showed that no multiplicative factor approximation exists. If the functions can be written as $f_i = g_i + h$ for some nonnegative monotone submodular g_i and a nonnegative symmetric submodular function h, a $O(\log(n))$ approximation is given in [3]. In the more general case where h is nonnegative submodular, a $O(k \cdot \log(n))$ approximation is provided in [7], and this is tight [35]. For nonnegative monotone functions, MSCA is equivalent to the Submodular Facility Location problem considered in [42], where a tight $O(\log(n))$ approximation is given.

Goel et al [12] consider the minimization case of (3) for nonnegative monotone submodular functions, in which $\mathcal{F}$ is a nontrivial collection of subsets of V (i.e. $\mathcal{F} \subset 2^V$) and there is no restriction on the $\mathcal{F}_i$ (i.e. $\mathcal{F}_i = 2^V$ for all i). In particular, given a graph G they consider the families of vertex covers, spanning trees, perfect matchings, and shortest st paths. They provide a tight $O(\log(n))$ approximation for the vertex cover problem, and show polynomial hardness for the other cases. To the best of our knowledge this is the first time that Problem (3) has been studied for nontrivial collections $\mathcal{F}$ of subsets of V.

2 Generic Multi-agent to Single-agent reduction

In this section, we describe a generic reduction of the capacitated multivariate problem (4) (and hence (3)) to a single-agent problem. The reduction transforms the solution space into a *lifted space*. The new space lives in an associated bipartite graph and facilitates the encoding of "ownership" in the multi-agent setting. For some special families $\mathcal{F}$, the lifted space has the same structure and hence the single and multi-agent problems behave similarly (w.r.t. approximability). The families of forests and matchings are examples of these. We call these families *invariant* and discuss them in Section 2.3.

2.1 The Lifting Reduction

In this section we describe a generic reduction of (3) or (4) to a problem of the form

$$\max / \min f(S) : S \in \mathcal{L}$$

for some function f and family of sets $\mathcal{L}$. This reduction, i.e. the definition of f and $\mathcal{L}$, is independent of whether we are considering a maximization or minimization problem. The argument is based on an observation of Nisan et al in [32], where they point out that the Submodular Welfare Problem (in the setting of combinatorial auctions) is a special case of submodular maximization subject to a matroid constraint.

Consider the complete bipartite graph $G = ([k]+V, E)$. We think of E as our new ground set (or lifted space), where an edge $(i, v) \in E$ corresponds to assigning item v to agent i.

Formally, we define a mapping π that identifies the set $(2^V)^k$ of possible tuples with the set 2^E of subsets of edges in the bipartite graph as follows, $\pi : (2^V)^k \to 2^E$ where

$$\pi(S_1, ..., S_k) = \biguplus_{i\in[k]} \{(i, v) \in E : v \in S_i\} = \biguplus_{i\in[k]} (\{i\} \times S_i).$$

Now for any $S \subseteq E$ and $i \in [k]$, let $S_i = \{v : (i, v) \in S\}$. One now easily checks that π is a bijection between the above two sets.

Claim. The mapping π is a bijection between $(2^V)^k$ and 2^E with $\pi^{-1}(S) = (S_1, S_2, \ldots, S_k)$.

Proof. Clearly π is one-to-one since every set $S \in 2^E$ (or equivalently $S \subseteq E$) is written uniquely as $S = \uplus_{i\in[k]}(\{i\} \times S_i)$. From this we see that π^{-1} is as claimed. □

We also see that union and intersection behaves nicely under the mapping π.

Claim. Let $S = \pi(S_1, ..., S_k)$ and $T = \pi(T_1, ..., T_k)$. Then

$$S \cup T = \pi(S_1 \cup T_1, S_2 \cup T_2, ..., S_k \cup T_k) \quad \text{and} \quad S \cap T = \pi(S_1 \cap T_1, S_2 \cap T_2, ..., S_k \cap T_k).$$

Proof. We show the union identity, the intersection argument being similar.

$$\begin{aligned} S \cup T &= \pi(S_1, ..., S_k) \cup \pi(T_1, ..., T_k) \\ &= (\uplus_{i\in[k]}\{(i, v) : v \in S_i\}) \cup (\uplus_{i\in[k]}\{(i, v) : v \in T_i\}) \\ &= \uplus_{i\in[k]}\{(i, v) : v \in S_i \cup T_i\} \\ &= \pi(S_1 \cup T_1, S_2 \cup T_2, ..., S_k \cup T_k). \end{aligned}$$

□

We can now go from a multivariate function g over V - such as the objectives in (3) and (4) - to a univariate function in the lifted space. Namely, one defines an associated univariate function $f : 2^E \to \mathbb{R}$ as follows:

$$f(S) = g(\pi^{-1}(S)),\ \forall S \subseteq E.$$

We next give a description of the feasible sets in the lifted space E. We define two families of sets over E that capture the information of the two constraints in the original class of problems:

$$\mathcal{H} := \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{F}\} \quad \text{and} \quad \mathcal{H}' := \{S \subseteq E : S_i \in \mathcal{F}_i,\ \forall i \in [k]\}.$$

Notice that in the definitions of $\mathcal{H}$ and $\mathcal{H}'$ we are implicitly using the fact that each $S \subseteq E$ can be written uniquely as $S = \uplus_{i\in[k]}(\{i\} \times S_i)$ where $S_i \subseteq V$.

Now, it should be clear that with the above definitions of $f, \mathcal{H}$ and $\mathcal{H}'$, we have:

$$\begin{array}{lll} \max / \min & g(S_1, ..., S_k) & \\ \text{s.t.} & S_1 + \cdots + S_k \in \mathcal{F} & \\ & S_i \in \mathcal{F}_i\,,\ \forall i \in [k] & \end{array} \quad = \quad \begin{array}{ll} \max / \min & f(S) \\ \text{s.t.} & S \in \mathcal{H} \cap \mathcal{H}'\,. \end{array}$$

Finally, by letting $\mathcal{L} := \mathcal{H} \cap \mathcal{H}'$, our original problem reduces to

$$\max / \min f(S) : S \in \mathcal{L}.$$

Notice that for the robust problem discussed in Theorem 3 we get

$$\max_{\substack{S_1+\cdots+S_k\in\mathcal{F}\\ S_i\in\mathcal{F}_i}} \min_{\substack{A_i\subseteq S_i\\ \sum_i|A_i|\le\tau}} g(S_1-A_1,\ldots,S_k-A_k) = \max_{S\in\mathcal{L}} \min_{\substack{A\subseteq S\\ |A|\le\tau}} f(S-A).$$

Clearly, this reduction is interesting if our new function f and family of sets $\mathcal{L}$ have properties which allows us to handle them computationally. These properties depend on the structure of the function g, and the families of sets $\mathcal{F}$ and $\mathcal{F}_i$'s. In the next section we discuss some properties that are preserved under this reduction.

2.2 Structure and properties preserved

In this section we show that several properties of the original function g and families $\mathcal{F}_i$ and $\mathcal{F}$ are preserved under the above reduction. For instance, if the original function g is (nonnegative, respectively monotone) multivariate submodular, then f is also (nonnegative, respectively monotone) submodular. Also, if the family $\mathcal{F}$ induces a matroid over the original ground set V, then so does $\mathcal{H}$ over the lifted space E. These results are summarized in the table below. Note that these, combined with the recent advances in SA submodular optimization (Section 1.4), now imply Theorem 2 and Theorem 3.

	MA	**SA**	Result
1	g multi-submodular	f submodular	Claim 2.2
2	g monotone	f monotone	Claim 2.2
3	$(V,\mathcal{F})$ a matroid	$(E,\mathcal{H})$ a matroid	Lemma 1
4	$(V,\mathcal{F})$ a p-matroid intersection	$(E,\mathcal{H})$ a p-matroid intersection	Corollary 2
5	$\mathcal{F}=$ bases of a matroid	$\mathcal{H}=$ bases of a matroid	Corollary 3
6	$\mathcal{F}=\{V\}$	$\mathcal{H}=$ set of bases of a partition matroid	Claim 2.2
7	$\mathcal{F}=2^V$	$(E,\mathcal{H})$ a partition matroid	Claim 2.2
8	$(V,\mathcal{F}_i)$ a matroid for all $i\in[k]$	$(E,\mathcal{H}')$ a matroid	Claim 2.2
9	$\mathcal{F}_i$ a ring family for all $i\in[k]$	$\mathcal{H}'$ a ring family	Claim 2.2

First, we consider which kind of properties our new function f inherits from the function g. The following two claims show that if g is a (nonnegative, respectively monotone) multivariate submodular function, then f as defined above is (nonnegative, respectively monotone) submodular.

Claim. The function f defined above is submodular if and only if g is multi-submodular.

Proof. By definition of f and Claim 2.1 we have that for any $S=\pi(S_1,...,S_k)$ and $T=\pi(T_1,...,T_k)$,

$$f(S)+f(T)\ge f(S\cup T)+f(S\cap T)$$
$$\Updownarrow$$
$$g(S_1,...,S_k)+g(T_1,...,T_k)\ge g(S_1\cup T_1,...,S_k\cup T_k)+g(S_1\cap T_1,...,S_k\cap T_k).$$

The result now follows from the fact that π is a bijection, that is, there is a one-to-one correspondence between sets $S\in 2^E$ and tuples $(S_1,...,S_k)\in(2^V)^k$. □

Claim. The function f is monotone if and only if g is monotone, and f is nonnegative if and only if g is nonnegative.

Proof. The latter statement follows directly from the definition of f. The monotonicity statement follows from noticing that given any two sets $S, T \subseteq E$, we have that $S \subseteq T$ if and only if the corresponding tuple of S is contained in the corresponding tuple of T, i.e., if and only if $(S_1, ..., S_k) \subseteq (T_1, ..., T_k)$. □

Next, we discuss some properties of $\mathcal{F}$ and the $\mathcal{F}_i$'s that are inherited by $\mathcal{H}$ and $\mathcal{H}'$ respectively.

Lemma 1. *If $(V, \mathcal{F})$ is a matroid, then $(E, \mathcal{H})$ is a matroid.*

Proof. We check that the two matroid axioms hold for $(E, \mathcal{H})$.

Hereditary property: If $T \in \mathcal{H}$ and $S \subseteq T$, then $S \in \mathcal{H}$. To see this, let $T \in \mathcal{H}$ and and assume $S \subseteq T$. We know that S and T can be written uniquely as $S = \uplus_{i\in[k]}(\{i\} \times S_i)$ and $T = \uplus_{i\in[k]}(\{i\} \times T_i)$ for some sets S_i's and T_i's in V. Since $S \subseteq T$, we must have that $S_i \subseteq T_i$ for each $i \in [k]$. Now, since $T \in \mathcal{H}$, we have that $T_i \cap T_j = \emptyset$ for all $i \neq j$. Hence, $S_i \cap S_j = \emptyset$ for all $i \neq j$. Also, since $T \in \mathcal{H}$, we have that $T_1 \cup ... \cup T_k \in \mathcal{F}$. And given that $S_1 \cup ... \cup S_k \subseteq T_1 \cup ... \cup T_k$ and $(V, \mathcal{F})$ is a matroid, it follows that $S_1 \cup ... \cup S_k \in \mathcal{F}$. Thus, $S \in \mathcal{H}$.

Greedy property: If $S, T \in \mathcal{H}$ and $|T| > |S|$, there exists $a \in T - S$ such that $S + a \in \mathcal{H}$. Again, let $S = \uplus_{i\in[k]}(\{i\} \times S_i)$ and $T = \uplus_{i\in[k]}(\{i\} \times T_i)$. Notice that $|T| = \sum_{i\in[k]} |T_i|$ and $|S| = \sum_{i\in[k]} |S_i|$. Let $I = S_1 \cup ... \cup S_k$ and $J = T_1 \cup ... \cup T_k$. Since $S, T \in \mathcal{H}$, we know that $I, J \in \mathcal{F}$. Also, by the disjoint condition on the sets S_i's and T_i's, we have that $|I| = \sum_{i\in[k]} |S_i|$ and $|J| = \sum_{i\in[k]} |T_i|$. Hence, $I, J \in \mathcal{F}$ and $|J| > |I|$. Since $(V, \mathcal{F})$ is a matroid, we know there must exist some $b \in J - I$ such that $I + b \in \mathcal{F}$. Now, since $b \in J$, there exists some $i_0 \in [k]$ such that $b \in T_{i_0}$, and hence $(i_0, b) \in T$. Also, since $b \notin I$, we have that $(i_0, b) \notin S$. Finally, we show that $S + (i_0, b) \in \mathcal{H}$. Observe that $S + (i_0, b) = [\uplus_{i\in[k]}(\{i\} \times S_i)] + (i_0, b) = [\uplus_{i\in[k]-i_0}(\{i\} \times S_i)] + (\{i_0\} \times S^*_{i_0})$ where $S^*_{i_0} = S_{i_0} + b$. We have to show that the sets $S_1, ..., S_{i_0-1}, S^*_{i_0}, S_{i_0+1}, ..., S_k$ are pairwise disjoint, and that their union is an independent set in $(V, \mathcal{F})$. But the second condition follows directly from the fact that $(\cup_{i\in[k]-i_0} S_i) \cup S^*_{i_0} = I + b \in \mathcal{F}$, while the first one from the fact that $b \notin I$. □

Corollary 2. *If $(V, \mathcal{F})$ is a p-matroid intersection, then $(E, \mathcal{H})$ is a p-matroid intersection.*

Proof. Let $\mathcal{F} = \cap_{i=1}^{p} \mathcal{I}_i$ for some matroids $(V, \mathcal{I}_i)$. Then we have that

$$\begin{aligned}
\mathcal{H} =& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{F}\} \\
=& \{S \subseteq E : S_1 + \cdots + S_k \in \cap_{i=1}^{p} \mathcal{I}_i\} \\
=& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{I}_i, \forall i \in [p]\} \\
=& \bigcap_{i\in[p]} \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{I}_i\} \\
=& \bigcap_{i\in[p]} \mathcal{I}'_i.
\end{aligned}$$

Moreover, from Lemma 1 we know that $(E, \mathcal{I}'_i)$ is a matroid for each $i \in [p]$, and the result follows. □

Corollary 3. *Assume $\mathcal{F}$ is the set of bases of some matroid $\mathcal{M} = (V, \mathcal{I})$, then $\mathcal{H}$ is the set of bases of some matroid $\mathcal{M}' = (E, \mathcal{I}')$.*

Proof. Let $\mathcal{F} = \mathcal{B}(\mathcal{M})$ and define $\mathcal{I}' = \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{I}\}$. From Lemma 1 we know that $\mathcal{M}' = (E, \mathcal{I}')$ is a matroid. We show that $\mathcal{H} = \mathcal{B}(\mathcal{M}')$. Let $r_{\mathcal{M}}$ and $r_{\mathcal{M}'}$ denote the rank of $\mathcal{M}$ and $\mathcal{M}'$ respectively. Then notice that

$$\begin{aligned}\mathcal{H} =& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{F}\}\\ =& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{B}(\mathcal{M})\}\\ =& \{S \subseteq E : S_1 + \cdots + S_k = I \in \mathcal{I}, |I| = r_{\mathcal{M}}\}\\ =& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{I}, |S| = r_{\mathcal{M}}\},\end{aligned}$$

where the last equality follows from the fact that $|I| = |S_1 + \cdots + S_k| = \sum_{i\in[k]} |S_i| = |S|$. On the other hand we have that

$$\begin{aligned}\mathcal{B}(\mathcal{M}') =& \{S \subseteq E : S \in \mathcal{I}', |S| = r_{\mathcal{M}'}\}\\ =& \{S \subseteq E : S_1 + \cdots + S_k \in \mathcal{I}, |S| = r_{\mathcal{M}'}\}.\end{aligned}$$

Hence, $\mathcal{H} = \mathcal{B}(\mathcal{M}')$ if and only if $r_{\mathcal{M}} = r_{\mathcal{M}'}$. But notice that for any $S \in \mathcal{I}'$ we have that $|S| = |I|$ for some $I \in \mathcal{I}$, and hence $r_{\mathcal{M}'} \leq r_{\mathcal{M}}$. Also, for any $I \in \mathcal{I}$ we can find some $S \in \mathcal{I}'$ (e.g. set $S_1 = I$ and $S_i = \emptyset$ for $i \neq 1$) such that $|S| = |I|$, and thus $r_{\mathcal{M}'} \geq r_{\mathcal{M}}$. It follows that $r_{\mathcal{M}} = r_{\mathcal{M}'}$ and hence $\mathcal{H} = \mathcal{B}(\mathcal{M}')$ as we wanted to show. □

Claim. Let $\mathcal{F} = 2^V$, then $\mathcal{H}$ correspond to the family of independent sets of a partition matroid.

Proof. Assume $\mathcal{F} = 2^V$, then

$$\begin{aligned}\mathcal{H} &= \{S \subseteq E : S_1 + \cdots + S_k \in 2^V\}\\ &= \{S \subseteq E : S_i \cap S_j = \emptyset,\ \forall i \neq j\}\\ &= \{S \subseteq E : \forall v \in V,\ |S \cap \delta(v)| \leq 1\}\\ &= \{S \subseteq E : \forall v \in V,\ |S \cap ([k] \times \{v\})| \leq 1\}.\end{aligned}$$

Since $\{[k] \times \{v\} : v \in V\}$ is a partition of E, it follows that $(E, \mathcal{H})$ is a partition matroid over E. □

Claim. If $\mathcal{F} = \{V\}$, then $\mathcal{H}$ corresponds to the set of bases of a partition matroid.

Proof. Let $\mathcal{F} = \{V\}$, then $\mathcal{H} = \{S \subseteq E : S_1 + \cdots + S_k = V\}$. From Claim 2.2 we know that $\mathcal{M} = (E, \mathcal{I})$ is a partition matroid over E, where

$$\mathcal{I} := \{S \subseteq E : \forall v \in V,\ |S \cap ([k] \times \{v\})| \leq 1\}.$$

Moreover, it is easy to see that the set of bases of $\mathcal{M}$ corresponds exactly to $\{S \subseteq E : S_1 + \cdots + S_k = V\}$. Since $\mathcal{H} = \{S \subseteq E : S_1 + \cdots + S_k = V\}$, the claim follows. □

Claim. If $(V, \mathcal{F}_i)$ is a matroid for each $i \in [k]$, then $(E, \mathcal{H}')$ is also a matroid.

Proof. Let $\mathcal{M}_i := (\{i\} \times V, \mathcal{I}_i)$ for $i \in [k]$, where $\mathcal{I}_i := \{\{i\} \times S : S \in \mathcal{F}_i\}$. Since $(V, \mathcal{F}_i)$ is a matroid, we have that $\mathcal{M}_i$ is also a matroid. Moreover, by taking the matroid union of the $\mathcal{M}_i$'s we get $(E, \mathcal{H}')$. Hence, $(E, \mathcal{H}')$ is a matroid. □

Claim. If $\mathcal{F}_i$ is a ring family over V for each $i \in [k]$, then $\mathcal{H}'$ is a ring family over E.

Proof. Let $S, T \in \mathcal{H}'$ and notice that $S \cup T = \biguplus_{i\in[k]}(\{i\} \times (S_i \cup T_i))$ and $S \cap T = \biguplus_{i\in[k]}(\{i\} \times (S_i \cap T_i))$. Since $\mathcal{F}_i$ is a ring family for each $i \in [k]$, it follows that $S_i \cup T_i \in \mathcal{F}_i$ and $S_i \cap T_i \in \mathcal{F}_i$ for each $i \in [k]$. Hence $S \cup T \in \mathcal{H}'$ and $S \cap T \in \mathcal{H}'$, and thus $\mathcal{H}'$ is a ring family over E. □

2.3 Invariant families

For a number of families $\mathcal{F}$ the approximations factors revealed in [12] are the same for both the single and multi-agent versions. The lifted space reduction from the previous section is useful for explaining why this occurs in several cases. In particular, for various families it turns out that the class $\mathcal{F}$ is preserved when performing this reduction. We call these *invariant families* (under the above MA to SA reduction). For instance, invariance in the case of forests could be stated as: if the multivariate problem is defined over the family $\mathcal{F}$ of forests of some graph G, then the reduced (single-agent) version can also be seen as defined over the family $\mathcal{F}'$ of forests of some graph G'. More formally, let E and $\mathcal{H}$ be as defined in the reduction. We say that the class of forests is invariant if: given any graph G and family $\mathcal{F}$ of forests of G, there exists a graph G' with set of edges E' and a bijection $\pi' : E \to E'$ satisfying that $\mathcal{F}' := \{\pi'(S) : S \in \mathcal{H}\}$ is the family of forests of G'. We show that forests (and spanning trees), matchings (and perfect matchings), and st-paths are examples of such families.

Lemma 2. *Given a graph G and a family of sets $\mathcal{F}$ over the edges of G, the multi-agent to single-agent reduction is invariant for the following classes:*

(1) Forests
(2) Matchings
(3) st paths

Proof. Let $E, \mathcal{H}$, and π be as defined in the reduction from Section 2. To be consistent with our previous notation we denote by V the set of edges of G, since this is the ground set of the original problem. Hence, in this case an element $v \in V$ denotes an edge of the original graph G. The lifting reduction is based on the idea of making k disjoint copies for each original element, and visualize the new ground set (or lifted space) as edges in a bipartite graph. However, when the original ground set corresponds to the set of edges of some graph G, we may just think of the lifted space in this case as being the set of edges of the graph G' obtained by taking k disjoint copies of each original edge. We think of choosing the edge that corresponds to the ith copy of v as assigning element v to agent i. We can formalize this by defining a mapping $\pi' : E \to E'$ that takes an edge $(i, v) \in E$ to the edge in G' that corresponds to the ith copy of edge v in G. It is clear that π' is a bijection. Then, via the bijections π and π' we can identify a tuple $(S_1, ..., S_k)$ with the set of edges in G' that takes the first copy of the edges in S_1, the second copies of the edges in S_2, and so on. Now, the following three statements are straightforward.

(1) Let $\mathcal{F}$ be the family of forests of G. Then: $(S_1, ..., S_k)$ is a feasible tuple for the original problem (i.e. $S_1 + \cdots + S_k \in \mathcal{F}$) $\iff$ $\pi(S_1, ..., S_k) \in \mathcal{H}$ $\iff$ $\pi'(\pi(S_1, ..., S_k))$ induces a forest in G'.

(2) Let $\mathcal{F}$ be the family of matchings in G. Then: $(S_1, ..., S_k)$ is a feasible tuple for the original problem $\iff$ $\pi(S_1, ..., S_k) \in \mathcal{H}$ $\iff$ $\pi'(\pi(S_1, ..., S_k))$ induces a matching in G'.

(3) Let $\mathcal{F}$ be the family of st-paths in G. Then: $(S_1, ..., S_k)$ is a feasible tuple for the original problem $\iff$ $\pi(S_1, ..., S_k) \in \mathcal{H}$ $\iff$ $\pi'(\pi(S_1, ..., S_k))$ induces an st-path in G'.

□

Notice that in the proof above $\pi' \circ \pi$ is a bijection, and thus the sizes of the sets are preserved. In particular, given a tuple $(S_1, ...S_k)$ such that $S_1 + \cdots + S_k = F \in \mathcal{F}$, we have that $|\pi'(\pi(S_1, ..., S_k))| = |\pi(S_1, ..., S_k)| = |F|$. This leads to the following corollary.

Corollary 4. *Given a graph G, both the families of spanning trees and perfect matchings of G are invariant under the lifting reduction.*

In addition, by using the bijection π' we can define a function $f' : 2^{E'} \to \mathbb{R}$ as $f'(S') = f(\pi'^{-1}(S'))$. It is clear that f' inherits all the properties (e.g. nonnegativity, monotonicity, submodularity, etc) of f. Then, we have the following:

$$\begin{array}{lll} \min & g(S_1, ..., S_k) & = \min \ f(S) \ = \min \ f'(S') \\ \text{s.t.} & S_1 + \cdots + S_k \in \mathcal{F} & \text{s.t. } S \in \mathcal{H} \quad \text{s.t. } S' \in \mathcal{F}', \end{array}$$

where $\mathcal{F}' = \{\pi'(S) : S \in \mathcal{H}\}$. Moreover, from Lemma 2 and Corollary 4 we have that $\mathcal{F}$ and $\mathcal{F}'$ are of the same class in the cases where $\mathcal{F}$ are forests, spanning trees, matchings, perfect matchings, and st-paths.

The following result describes the algorithmic consequences of Lemma 2 and Corollary 4.

Corollary 5. *The SA and MV problems for the structures described in Lemma 2 and Corollary 4 are equivalent, in the sense that approximation factors and hardness results are the same for the SA and MV versions of the problem.*

Proof. It follows from the fact that approximation factors and hardness results for the above combinatorial structures only depend on the number of nodes (and not on the number of edges) of the underlying graph. Given that the number of nodes remain invariant under taking k disjoint copies of each edge, the result follows. □

2.4 Invariance of multivariate rings

In this section we focus on the class of problems

$$\begin{array}{ll} \min / \max & g(S_1, ..., S_k) \\ \text{s.t.} & (S_1, ..., S_k) \in \mathcal{D}, \end{array}$$

where $\mathcal{D} \subseteq (2^V)^k$. Notice that this is a much more general framework than the CMVSO class of problems defined in (4). In general, one would not expect to have much tractability for these kind of problems even when both the objective function and the feasible space have some structure. In this section we show that the generic reduction from Section 2 allows us to solve the above problem exactly in the minimization setting when the objective function g is an arbitrary multivariate submodular function and the feasible space $\mathcal{D}$ has a ring like structure (that is, closed under unions and intersections). Formally, we say that a family of tuples $\mathcal{D}$ is a *multivariate ring* if for any two tuples $(S_1, ..., S_k), (T_1, ..., T_k) \in \mathcal{D}$ the coordinate-wise union and intersection (i.e. $(S_1 \cup T_1, S_2 \cup T_2, ..., S_k \cup T_k)$ and $(S_1 \cap T_1, ..., S_k \cap T_k)$) also belong to $\mathcal{D}$. The following result shows that the ring like structure is preserved under the reduction.

Lemma 3. *Let E and π be as defined in the reduction of Section 2. Then, if $\mathcal{D}$ is a multivariate ring, the family $\mathcal{D}' := \{\pi(S_1, ..., S_k) : (S_1, ..., S_k) \in \mathcal{D}\}$ is a ring over E.*

Proof. Let $S, T \in \mathcal{D}'$. Let $(S_1, ..., S_k), (T_1, ..., T_k) \in \mathcal{D}$ be such that $S = \pi(S_1, ..., S_k)$ and $T = \pi(T_1, ..., T_k)$. We have to show $S \cup T, S \cap T \in \mathcal{D}'$. From Claim 2.1 we know that $S \cup T = \pi(S_1 \cup T_1, S_2 \cup T_2, ..., S_k \cup T_k)$ and $S \cap T = \pi(S_1 \cap T_1, ..., S_k \cap T_k)$. Now the result follows from the fact that $(S_1 \cup T_1, ..., S_k \cup T_k), (S_1 \cap T_1, ..., S_k \cap T_k) \in \mathcal{D}$ by definition of multivariate ring. □

Now, it follows from the reduction that

$$\begin{array}{ll} \min & g(S_1, ..., S_k) \\ \text{s.t.} & (S_1, ..., S_k) \in \mathcal{D} \end{array} = \begin{array}{ll} \min & f(S) \\ \text{s.t.} & S \in \mathcal{D}', \end{array}$$

where if $\mathcal{D}$ is multivariate ring then $\mathcal{D}'$ is a ring, and if g is an arbitrary multivariate submodular function then f is an arbitrary submodular function. Hence, given that minimization of a submodular function over a ring family can be done efficiently ([38,40]), we have the following result.

Lemma 4. *The problem*

$$\min g(S_1, ..., S_k) : (S_1, ..., S_k) \in \mathcal{D}$$

where $\mathcal{D}$ is a multivariate ring family and g is an arbitrary multivariate submodular function can be solved exactly in polynomial time.

3 Reductions for Multi-Agent and Multivariate Minimization

In this section we seek a generic reduction of a multi-agent (and some times multivariate) minimization problem to its single-agent primitive. We develop this method in the case where the objective function is nonnegative monotone (multivariate) submodular and we work with the natural LP formulation (cf. Appendix A). Specifically we prove the following.

Theorem 6. *Suppose there is a (polytime) $\alpha(n)$-approximation for the single-agent nonnegative monotone submodular minimization problem based on rounding the natural LP for a family $\mathcal{F}$. Then there is a (polytime) $O(\alpha(n) \cdot \min\{k, \log^2(n)\})$-approximation for the multi-agent nonnegative monotone submodular minimization problem over $\mathcal{F}$. Moreover, the $O(k\alpha(n))$-approximation extends to the more general multivariate setting.*

We note that the $\log^2(n)$ approximation loss due to having multiple agents is in the right ballpark. This is because for the vertex cover problem there is a factor 2-approximation for single-agent submodular minimization, and a (tight) $O(\log(n))$-approximation for the multi-agent version [12].

3.1 The Single-Agent and Multi-Agent Formulations

Due to monotonicity, one may often assume that we are working with a family $\mathcal{F}$ which is *upwards-closed*, aka a *blocking family*. The advantage is that to certify whether $F \in \mathcal{F}$, we only need to check that $F \cap B \neq \emptyset$ for each element B of the family $\mathcal{B}(\mathcal{F})$ of minimal blockers of $\mathcal{F}$. We discuss the details in Appendix A. This gives rise to the natural blocking LP formulation for the single-agent and multi-agent problems as used in [12]. These are discussed in detail in Appendix B. The *single-agent LP* is:

$$\begin{array}{ll} \min & \sum_{S \subseteq V} f(S)x(S) \\ \text{s.t.} & \sum_{v \in B} \sum_{S \ni v} x(S) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F}) \\ & x(S) \geq 0, \end{array} \tag{5}$$

and the corresponding dual is given by

$$\begin{array}{ll} \max & \sum_{B\in\mathcal{B}(\mathcal{F})} y_B \\ \text{s.t.} & \sum_{v\in S}\sum_{B\ni v} y_B \le f(S),\ \forall S\subseteq V \\ & y_B \ge 0. \end{array} \tag{6}$$

The *multi-agent LP* is:

$$\begin{array}{ll} \min & \sum_{S\subseteq V}\sum_{i\in[k]} f_i(S)x(S,i) \\ \text{s.t.} & \sum_{v\in B}\sum_{S\ni v}\sum_{i\in[k]} x(S,i) \ge 1,\ \forall B\in\mathcal{B}(\mathcal{F}) \\ & x(S,i)\ge 0, \end{array} \tag{7}$$

and its dual is:

$$\begin{array}{ll} \max & \sum_{B\in\mathcal{B}(\mathcal{F})} y_B \\ \text{s.t.} & \sum_{v\in S}\sum_{B\ni v} y_B \le f_i(S),\ \forall i\in[k],\ \forall S\subseteq V \\ & y_B \ge 0. \end{array} \tag{8}$$

By standard methods (see Appendix B) one may solve these problems exactly in polytime if one can separate over the blocking formulation $P^*(\mathcal{F}) := \{z \ge 0 : z(B) \ge 1 \textit{ for all } B \in \mathcal{B}(\mathcal{F})\}$. This is often the case for many natural families such as spanning trees, perfect matchings, *st*-paths, and vertex covers. For any feasible solution x to (7) we associate a vector $z : V \to \mathbf{R}_+$ by $z(v) := \sum_{S\ni v}\sum_{i\in[k]} x(S,i)$. Note that feasibility of x is equivalent to the condition $z \in P^*(\mathcal{F})$. In the following we work with the extended vectors (x, z).

3.2 A Multi-Agent $O(\alpha(n)\log^2(n)) - approximation$

While the proof of this result uses several standard pre-processing ingredients, it also requires some new techniques to deal with submodular objectives (these would not be necessary for a multi-agent problem with linear (modular) objectives). We also need the submodularity of a particular set function which arises from a partition $U_1, U_2, \ldots, U_k$ of a ground set V. Suppose we are given such sets with arbitrary costs $c(U_i)$. We define a set function $g : 2^V \to \mathbb{R}$ where $g(\emptyset) = 0$ and for any non-empty set S:

$$g(S) = \sum_{i:S\cap U_i\neq\emptyset} c(U_i).$$

Proposition 2. *The function $g(S)$ is submodular and monotone if the $c(U_i)$ are nonnegative.*

Proof. We establish the diminishing returns property of g. Let $A \subseteq B$ and $v \in V - B$. Suppose that $g(B + v) - g(B) > 0$. This means that $v \in U_i$ such that $B \cap U_i = \emptyset$. But then $A \cap U_i = \emptyset$ as well, and so $g(A + v) - g(A) = c(U_i) = g(B + v) - g(B)$. □

Theorem 7. *Suppose there is a (polytime) $\alpha(n)$-approximation for the single-agent nonnegative monotone submodular minimization problem based on rounding the natural LP for a family $\mathcal{F}$. Then there is a (polytime) $O(\alpha(n)\log^2(n))$-approximation for the multi-agent nonnegative monotone submodular minimization problem over $\mathcal{F}$.*

Proof. Let (x^*, z^*) denote an optimal solution to the multi-agent LP associated with $\mathcal{F}$ and the f_i's. As discussed in Appendix B.4, this can be found in polytime under fairly mild assumptions. In order to apply a black box single-agent rounding algorithm we must create

a different multi-agent solution. This is done in several steps, the first few of which are standard. The key steps are the *fracture, expand and return* steps which arise later in the process.

Call an element v *small* if $z^*(v) \leq \frac{1}{2n}$. Note that $\sum_{v \text{ small}} z^*(v) \leq \frac{1}{2}$, and so for any blocking set $B \in \mathcal{B}(\mathcal{F})$, we have that at most $\frac{1}{2}$ of $z^*(B)$ is contributed by small elements. Let S^- denote the set obtained after removing all small elements from S. Then for any $B \in \mathcal{B}(\mathcal{F})$ we have $z^*(B^-) \geq \frac{1}{2}$. So let us consider dropping all small v by removing them from sets in the support and then doubling the resulting sets. That is, we obtain a new solution (x', z') by starting with $x' \equiv x^*$ and then doing the following for each (S, i) in the support of x^*. Reduce $x'(S, i)$ by $x^*(S, i)$ and then increase $x'(S^-, i)$ by the amount $2x^*(S, i)$. As f_i's are monotone, this at most doubles the cost of the original solution (x^*, z^*).

We now prune the solution (x', z') a bit more. Let Z_j be the elements v such that $z'(v) \in (2^{-(j+1)}, 2^{-j}]$ for $j = 0, 1, 2, \ldots, L$. Since $z'(v) > \frac{1}{2n}$ for any element in the support, we have that $L \leq \log(n)$. We call Z_j *bin* j and define $r_j = 2^j$. We round up each $v \in Z_j$ so that $z'(v) = 2^{-j}$ by augmenting the x' values by at most a factor of 2. We may do this simultaneously for all v by possibly "truncating" some of the sets involved, i.e., changing some $x'(S, i)$ values to be $x'(T, i)$ where $T \subseteq S$. As before, this is fine since f_i's are monotone. In the end, we call this a *uniform solution* (x'', z'') in the sense that each $z''(v)$ is some power of 2. Note that its cost is at most a constant factor times the original MA LP solution.

Fracture. We now create another *fractured* solution (x''', z''') as follows. For each support variable $x''(S, i) > 0$, we replace this variable by several variables, one for each set $S_j = S \cap Z_j$ with $j = 0, 1, 2, \ldots L$. Set $x'''(S_j, i) = x''(S, i)$ for each j. By monotonicity we have that $f_i(S_j) \leq f_i(S)$ and so the cost of this solution goes up by a factor of at most $L \leq \log(n)$. Note that $z''' = z''$.

Expand. Now for any $S \subseteq Z_j$ with $x'''(S, i) > 0$ for some i, we blow it up by a factor r_j. (Don't worry, this scaling is temporary.) Since $z'''(v) = \frac{1}{r_j}$ for each $v \in Z_j$, this means that the resulting values yield a (probably fractional) cover of Z_j. At this point, we are completely ignoring submodularity - we simply have sets with costs: e.g., for set S above we have $c(S) = f_i(S)$. By standard greedy, we then turn this fractional cover into an integral cover of Z_j by increasing the cost by a factor of at most $\ln(n)$. Moreover, we may again use monotonicity to truncate some of the sets so that our cover consists of disjoint sets. So let $U_{1j}, U_{2j}, \ldots U_{n(j)j}$ be the corresponding partition of Z_j.

Return. Now we go back to get a new MA LP solution $(\hat{x}, \hat{z})$ whose support is on disjoint sets. For each U_{sj} we re-set $\hat{x}(U_{sj}, i) = \frac{1}{r_j}$ where i was the agent associated with the set U_{sj}. Note that $\hat{z} = z'''$ and so this is indeed feasible (and again uniform).

Single-Agent Rounding. In the last step, we show that we may use Proposition 2 to apply single-agent submodular rounding to produce a provably "good" solution for the multi-agent problem.

We define the submodular function g on V induced by the disjoint sets U_{sj} and their costs which are inherited from various agents. We may then consider an associated single-agent submodular optimization problem $\min g(S) : S \in \mathcal{F}$. Clearly the multi-agent solution $(\hat{x}, \hat{z})$ induces a solution $(y, \hat{z})$ for this single-agent problem as well. Namely y has a support $\{U_{sj}\}$ and $y(U_{sj}) = \hat{x}(U_{sj}, i)$, where i is the agent associated to U_{sj}. Moreover, by definition of $g()$, the cost of this solution is the same as for $(\hat{x}, \hat{z})$.

Now suppose we have an α-approximation for submodular optimization over the natural formulation $P^*(\mathcal{F})$, and so we can produce an integral solution whose cost is at most α times the cost of $(y, \hat{z})$ which in turn costs at most $O(\log^2(n))$ MA-LP. (We paid a factor $O(\log n)$ for fracturing, and a factor $O(\ln n)$ for greedy set cover.) But an integral solution

for this single-agent problem is effectively picking off some of the sets U_{sj}'s, and the price we pay for any such set is the same as the associated agent's cost in the original MA optimization problem. If some agent i ends up picking several U_{sj}'s we may combine them due to subaddivity of f_i. In other words, we now have a feasible solution for the multi-agent problem whose cost is at most $O(\alpha log^2(n))$ times the original LP. This completes the description of the algorithm. □

3.3 A Multivariate $O(k\alpha(n)) - approximation$

In this section we consider the more general multivariate class of problems

$$\text{(MV)} \quad \min g(S_1, \ldots, S_k) : S_1 + \cdots + S_k \in \mathcal{F}, \tag{9}$$

where g is a nonnegative monotone multivariate submodular function. As shown in Proposition 1, this class encodes much more than their decomposable counterparts.

We use the generic reduction from Section 2 to obtain a *Single-agent Reduced* (SAR) problem

$$\text{(SAR)} \quad \min f(S) : S \in \mathcal{H}. \tag{10}$$

Let $G = ([k] + V, E)$ denote the underlying bipartite graph. Recall that $S \subseteq E(G)$ lies in $\mathcal{H}$ if and only if $cov(S) \in \mathcal{F}$ and $|S \cap \delta_G(v)| \leq 1 \ \forall v \in V$. Recall also that if g is nonnegative monotone multivariate submodular, then f is nonnegative monotone submodular (see Claims 2.2 and 2.2).

We first show that under mild assumptions we can solve in polytime the single-agent LP associated to (10).

Claim. If we can separate in polytime over $P^*(\mathcal{F})$, then we can solve in polytime the single-agent LP associated to (10).

Proof. First notice that the blocking sets of $\mathcal{H}$ are just $\mathcal{B}(\mathcal{H}) = \{\cup_{v \in B} \delta_G(v) : B \in \mathcal{B}(\mathcal{F})\}$. Then observe that

$$\begin{aligned} P^*(\mathcal{H}) &= \{w \in \mathbb{R}_+^E : w(B') \geq 1, \ \forall B' \in \mathcal{B}(\mathcal{H})\} \\ &= \{w \in \mathbb{R}_+^E : w(\cup_{v \in B} \delta_G(v)) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F})\} \\ &= \{w \in \mathbb{R}_+^E : \sum_{v \in B} w(\delta_G(v)) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F})\}. \end{aligned}$$

Now, given any vector $w \in \mathbb{R}_+^E$ we can associate a vector $z \in \mathbb{R}_+^V$ defined by $z(v) = w(\delta_G(v))$. Then, it is clear that for each $B \in \mathcal{B}(\mathcal{F})$ we have

$$z(B) \geq 1 \iff \sum_{v \in B} w(\delta_G(v)) \geq 1.$$

Hence, $w \in P^*(\mathcal{H})$ if and only if $z \in P^*(\mathcal{F})$. Given that we can separate in polytime over $P^*(\mathcal{F})$, it follows that we can also separate in polytime over $P^*(\mathcal{H})$. Hence we can solve the LP associated to (10) in polytime. □

We now prove the main result of this section. The proof combines the machinery from the reduction of Section 2 and some of the methods from Section 3.1.

Theorem 8. *Suppose there is a (polytime) $\alpha(n)$-approximation for the single-agent nonnegative monotone submodular minimization problem based on rounding the natural LP for a family $\mathcal{F}$. Then there is a (polytime) $O(k\alpha(n))$-approximation for the multivariate nonnegative monotone submodular minimization problem over $\mathcal{F}$.*

Proof. Consider the multivariate problem over $\mathcal{F}$ given by (9) and apply the generic reduction from Section 2 to get a single-agent problem over the lifted space as given by (10). Let OPT(MV) denote the value of the optimal solution to (9). Also, let OPT(SAR) denote the value of the optimal solution to (10), and OPT(SAR-LP) the value of the optimal solution to the LP associated with (10). Since (9) and (10) are equivalent, it is clear that

$$OPT(\text{SAR-LP}) \leq OPT(\text{SAR}) = OPT(\text{MV}).$$

We construct a new single-agent minimization problem:

$$(\text{NEW-SA}) \quad \min f'(S) : S \in \mathcal{F}$$

where $f' : 2^V \to \mathbb{R}_+$ is nonnegative monotone submodular. It's blocking LP relaxation also satisfies $OPT(\text{NEW-LP}) \leq k \cdot OPT(\text{MV})$. In addition, each feasible solution $F \in \mathcal{F}$ for (NEW-SA) induces a feasible solution $S_1 + \cdots + S_k = F$ for the original multivariate problem (MV) of the same objective value. Then the desired result follows from the (polytime) $\alpha(n)$-approximation assumption, since then we can find $F^* \in \mathcal{F}$ such that $F^* = S_1^* + \cdots + S_k^*$ and

$$g(S_1^*, \ldots, S_k^*) = f'(F^*) \leq \alpha \cdot OPT(\text{NEW-LP}) \leq \alpha k \cdot OPT(\text{MV}).$$

We construct the function f' as follows. Let $y^*(S)$ be an optimal solution for the single-agent LP associated to (10). Since y^* is feasible for the LP, we must have

$$\sum_{e \in B'} \sum_{S \ni e} y^*(S) \geq 1, \ \forall B' \in \mathcal{B}(\mathcal{H}).$$

Equivalently, since $\mathcal{B}(\mathcal{H}) = \{\cup_{v \in B} \delta_G(v) : B \in \mathcal{B}(\mathcal{F})\}$, the feasibility condition can be also stated as

$$\sum_{v \in B} \sum_{e \in \delta_G(v)} \sum_{S \ni e} y^*(S) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F}).$$

For any feasible solution y to (10) we associate a vector $w : E \to \mathbb{R}_+$ by $w(e) = \sum_{S \ni e} y(S)$. Note that feasibility of y is equivalent to the condition $w \in P^*(\mathcal{H})$. Let w^* be the vector associated to the optimal solution y^*. Moreover, for each $v_j \in V$ let e_j^* be the edge in the bipartite graph G satisfying

$$w^*(e_j^*) = \max_{e \in \delta_G(v_j)} w^*(e) = \max_{i \in [k]} w^*(i, v_j),$$

and let $T^* = \{e_1^*, e_2^*, \ldots, e_n^*\}$. Also, for any set $S \subseteq E$, let S^- denote the set obtained after removing all elements of $E - T^*$ from S.

So let us consider dropping all elements from $E - T^*$ by removing them from sets in the support and then increasing by a factor of k the resulting sets. That is, we obtain a new solution y' by starting with $y' \equiv y^*$ and then doing the following for each S in the support of y^*. Reduce $y'(S)$ by $y^*(S)$ and then increase $y'(S^-)$ by the amount $ky^*(S)$. As

f is monotone, this increases the cost of the original solution y^* by at most a factor of k. Notice that feasibility of y' follows by observing that for each $B \in \mathcal{B}(\mathcal{F})$ we have that

$$\sum_{v_j \in B} \sum_{e \in \delta(v_j)} \sum_{S \ni e} y'(S) = \sum_{v_j \in B} \sum_{S \ni e_j^*} y'(S) = \sum_{v_j \in B} \sum_{S \ni e_j^*} k y^*(S) \geq \sum_{v_j \in B} \sum_{e \in \delta(v_j)} \sum_{S \ni e} y^*(S) \geq 1.$$

Let $\pi : V \to T^*$ be the natural bijection between V and T^* where $\pi(v_j) = e_j^*$. Notice that for any set $S \subseteq V$ we have that $cov(\pi(S)) = S$ and $|\pi(S) \cap \delta_G(v)| \leq 1 \ \forall v \in V$. Hence

$$F \in \mathcal{F} \iff \pi(F) \in \mathcal{H}|_{T^*}.$$

Thus, any feasible solution $F \in \mathcal{F}$ induces (via the bijection π) a feasible solution for (SAR), which in turn induces (via the generic reduction from Section 2) a feasible solution for (MV).

Consider the function $f' : 2^V \to \mathbb{R}_+$ defined as $f'(S) = f(\pi(S))$. The right hand side is just the restriction $f|_{T^*}$ of f to T^* and hence it is again nonnegative monotone submodular. Since π is a bijection, f' is as well. Moreover, notice that for any $F \in \mathcal{F}$ we have (by definition of f') that $f'(F) = f(\pi(F))$, and hence

$$\min f'(S) : S \in \mathcal{F} \ = \ \min f(S) : S \in \mathcal{H}|_{T^*}.$$

In addition, $F^* \in \mathcal{F}$ is an optimal solution for $\min f'(S) : S \in \mathcal{F}$ if and only if $\pi(F^*)$ is an optimal solution for $\min f(S) : S \in \mathcal{H}|_{T^*}$.

We now consider the new single-agent problem $\min f'(S) : S \in \mathcal{F}$ and its associated LP given by:

$$\text{(NEW-LP)} \quad \begin{array}{ll} \min & \sum_{S \subseteq V} f'(S) x(S) \\ \text{s.t.} & \sum_{v \in B} \sum_{S \ni v} x(S) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F}) \\ & x(S) \geq 0. \end{array} \tag{11}$$

We construct a feasible solution x' for (11) by setting $x'(S) = y'(\pi(S))$ for each $S \subseteq V$. Or equivalently, for each $S \subseteq E$ in the support of y' set $x'(\pi^{-1}(S)) = y'(S)$, and set $x'(\pi^{-1}(S)) = 0$ for each $S \subseteq E$ not in the support. Notice that feasibility of x' for (11) now comes from the feasibility of y' for the LP associated to (10), since for any $B \in \mathcal{B}(\mathcal{F})$ we have

$$\sum_{v_j \in B} \sum_{S \ni v_j} x'(S) = \sum_{v_j \in B} \sum_{S \ni v_j} y'(\pi(S)) = \sum_{v_j \in B} \sum_{S \ni e_j^*} y'(S) \geq 1.$$

Moreover, we have that the objective value associated to the solution x' satisfies

$$\sum_{S \subseteq V} f'(S) x'(S) = \sum_{S \subseteq V} f(\pi(S)) y'(\pi(S)) = \sum_{S \subseteq E} f(S) y'(S) \leq k \cdot OPT(\text{SAR-LP}),$$

and hence

$$OPT(\text{NEW-LP}) \leq k \cdot OPT(\text{SAR-LP}) \leq k \cdot OPT(\text{MV})$$

as we wanted to show.

□

The above theorem has many interesting consequences. We now discuss one that leads to a polytime k-approximation for the submodular facility location problem, where k denotes the number of facilities.

Corollary 6. *Let f be a nonnegative monotone multivariate submodular function on V. Then there is a polytime k-approximation for*

$$\begin{aligned} \min \quad & f(S_1, ..., S_k) \\ s.t. \ & S_1 + \cdots + S_k = V. \end{aligned}$$

Proof. Notice that the single-agent version of the above multivariate problem is the trivial $\min f(S) : S \in \{V\}$. Hence a polytime exact algorithm is available for the single-agent problem and thus a polytime k-approximation is available for the multivariate version. □

A very special case of Corollary 6 occurs when $f(S_1, ..., S_k) = \sum_{i \in [k]} f_i(S_i)$ for some nonnegative monotone submodular f_i's. This is equivalent to the submodular facility location problem considered by Sviktina and Tardos in [42], where they give a tight $O(\ln(n))$ approximation. Their approximation is in terms of n, which in this setting denotes the number of customers/clients/demands. Notice that from Corollary 6 we can immediately provide an approximation in terms of the number of facilities. This bound clearly becomes preferable in facility location problems (for instance, Amazon) where the number of customers swamps the number of facility locations (i.e. $n >> k$).

Corollary 7. *There is a polytime k-approximation for submodular facility location, where k denotes the number of facilities.*

4 Bounded Blockers and Submodular Minimization

Many natural classes $\mathcal{F}$ have been shown to behave poorly for (nonnegative, monotone) submodular minimization. For instance, polynomial inapproximability has been established in the case where $\mathcal{F}$ is the family of edge covers ([19]), spanning trees ([12]), or perhaps most damning is $\tilde{\Omega}(\sqrt{n})$-inapproximability for minimizing $f(S)$ subject to the simple cardinality constraint $|S| \geq k$[4] [41]. These results all show that it is difficult to distinguish between two similar submodular functions, one of which "hides" the optimal set. At some level, the proofs leverage the fact that f can have wild (so-called) curvature [6].

A few cases exist where good approximations hold. For instance, for ring families the problem can be solved exactly by adapting any algorithm that works for unconstrained submodular minimization ([38,40]). Grötschel, Lovász, and Schrijver ([15]) show that the problem can be also solved exactly when $\mathcal{F}$ is a *triple family*. Examples of these kind of families include the family of all sets of odd cardinality, the family of all sets of even cardinality, or the family of sets having odd intersection with a fixed $T \subseteq V$. More generally, Goemans and Ramakrishnan ([13]) extend this last result by showing that the problem can be solved exactly when $\mathcal{F}$ is a *parity family* (a more general class of families). In the context of NP-Hard problems, the cases in which a good (say $O(1)$ or $O(\log(n))$) approximation exists are almost non-existent. We have that the submodular vertex cover admits a 2-approximation ([12,19]), and the k-uniform hitting set has $O(k)$-approximation. It is natural to examine the core reasons why these approximations algorithms work in the hope of harnessing ingredients in other contexts of interest.

We emphasize that our work in this section focuses only on nonnegative normalized monotone functions. We also restrict attention to upwards-closed families $\mathcal{F}$, that is, if $F \subseteq F'$ and $F \in \mathcal{F}$, then $F' \in \mathcal{F}$. Due to monotonicity this is usually a benign assumption

[4] which is equivalent to the constraint $|S| = k$ since f is monotone

(see Appendix A). For instance, if we have a well-described formulation (or approximation) for the polytope $P(\mathcal{F})$ whose vertices are $\{\chi^F : F \in \mathcal{F}\}$, then we can also separate over its upwards closure $P(\mathcal{F})^\uparrow := \{z \geq x : x \in P\}$. A second issue one must address when working with the upwards-closure of a family $\mathcal{F}$ is whether, given F' in the closure, one may find a set $F \in \mathcal{F}$ with $F \subseteq F'$ in polytime. This is also the case if a polytime separation oracle is available for $P(\mathcal{F})$.

4.1 Bounded Blockers and Single-Agent Minimization

Algorithms for vertex cover and k-uniform hitting sets rely only on the fact that the feasible space has a *bounded blocker property*. We call a clutter (set family) $\mathcal{F}$ *β-bounded* if $|F| \leq \beta$ for all $F \in \mathcal{F}$.

The next result is the main algorithmic device of this section and has several natural applications including implications for the class of multi-agent problems over $\mathcal{F}$. These are discussed in the following sections.

Theorem 9. *Consider the (nonnegative, normalized, monotone) submodular minimization primitive*

$$\min f(S) : S \in \mathcal{F}$$

and its natural LP (5). If $\mathcal{B}(\mathcal{F})$ is β-bounded, then (5) has integrality gap of at most β.

Proof. Now assume we have that x^* is an optimal solution for the primal (5) with value OPT. We claim that $Q = \{v \in V : \sum_{S:S\ni v} x^*(S) \geq \frac{1}{\beta}\}$ is an β-approximation. Recall that $Q \in \mathcal{F}$ if and only if $Q \cap B \neq \emptyset$ for each $B \in \mathcal{B}(\mathcal{F})$. Pick an arbitrary $B \in \mathcal{B}(\mathcal{F})$, since x^* is a feasible solution we must have that $\sum_{v\in B} \sum_{S\ni v} x^*(S) \geq 1$. Hence, there must exist some $v_0 \in B$ such that $\sum_{S\ni v_0} x^*(S) \geq \frac{1}{|B|} \geq \frac{1}{\beta}$, where the last inequality follows from the fact that $\mathcal{B}(\mathcal{F})$ is β-bounded. It follows that $v_0 \in Q$ and thus $Q \cap B \supseteq \{v_0\} \neq \emptyset$. Since βx^* is a fractional cover of Q, we may now apply Corollary 13 (the latter part in particular) to deduce that $f(Q) \leq \beta$ OPT. □

In many circumstances this result becomes algorithmic. For instance, as discussed in Lemma 11 in Appendix B.2 we have the following.

Corollary 8. *If $\mathcal{B}(\mathcal{F})$ is β-bounded and $|\mathcal{B}(\mathcal{F})| \in poly(n)$, then there is a polytime β-approximation for (5). This is the case in particular if $\beta = O(1)$.*

We note that it is not necessary to have a $O(1)$-bounded blocker to derive an $O(1)$-approximation algorithm; for instance we could have that $\mathcal{F}$ is polynomially bounded in some cases so an exact algorithm exists. The inapproximability result of Khot et al. [22] could be seen as a type of converse to Theorem 9; it suggests that if the blockers are large, and suitably rich, then submodular minimization over $\mathcal{F}$ is doomed.

Theorem 10 (Khot, Regev [22]). *Assuming the Unique Games conjecture, there is a $(k - \epsilon)$-inapproximability factor for the problem of finding a minimum size vertex cover in a k-uniform hypergraph.*

The restriction that $|\mathcal{B}(\mathcal{F})| \in poly(n)$ could be quite severe. In Appendix B.3 we discuss how LP (5) can still be solved in polytime under a much weaker assumption. In particular we show the following.

Corollary 9. *Assume there is a polytime separation oracle for the blocking formulation $P^*(\mathcal{F})$. Then LP (5) can be solved in polytime.*

4.2 Applications of Bounded Blockers

We discuss several problems which can be addressed by bounded blockers. We start with some of the simplest examples which also indicate how one might find uses for Theorem 9. We later discuss applications to multi-agent minimization.

*Example 8 (***Edge Covers).** *Let $G = (V, E)$ be an undirected graph and $\mathcal{F}$ denote the family of edge covers of G. It follows that $\mathcal{B}(\mathcal{F})$ consists of the n sets $\delta(v)$ for each $v \in V$. Theorem 9 immediately implies that that the minimum submodular edge cover problem is Δ-approximable for graphs of max degree Δ. In particular, the $\Omega(n/log(n))$ inapproximability result from [19] is pinpointed on graphs with at least one high-degree node.*

This example is not surprising but gives clarity to the message that whenever the blocker $\mathcal{B}(\mathcal{F})$ induces a bounded-degree hypergraph, Theorem 9 applies. Let's explore a few questions which can be approached like this. In [42] an $O(\ln n)$-approximation is given for submodular facility location, but for instances where the number of customers n dwarfs the number of facilities k the following is relevant.

*Example 9 (***Facility Location).** *Suppose we have a submodular facility location problem with k facilities and n customers. We may represent solutions as subsets of edges in a $k \times n$ bipartite graph G which include exactly one edge incident to each customer node. As our submodular functions are monotone, we can work with the associated blocking family $\mathcal{F}$ for which $\mathcal{B}(\mathcal{F})$ consists of the edges $\delta_G(v)$, where v is a customer node. A k-approximation now follows from Theorem 9.*

We consider the several examples from the broad task of finding *low-cost dense networks*.

*Example 10 (***Lightly Pruned Networks).** *Suppose we seek low-cost subgraphs of some graph G where nodes can only have a limited number τ of their edges removed. In other words, the degree of each node v should remain at least $deg(v) - \tau$. The family of subgraphs (viewed as edge subsets) $\mathcal{F}$ has as its blocker the set of all sub-stars of size $\tau + 1$: $\mathcal{B}(\mathcal{F}) = \{E' \subseteq \delta(v) : |E'| = \tau + 1, v \in V(G)\}$. When τ is small (constant or polylog) one obtains sensible approximations. This can be extended to notions of density which incorporate some tractable family of cuts $\delta(S) \in \mathcal{C}$. That is, where the designed subgraph must retain all but a bounded number τ of edges from each $\delta(S)$.*

*Example 11 (***Girth-Widening).** *For an undirected graph G we seek to minimize $f(J)$ where J is subset of edges for which $G - J$ does not have any short cycles. For instance, given a threshold τ we want $G - J$ to not have any cycles of length at most τ. We seek $J \in \mathcal{F}$ whose blocker consists of all cycles with at most τ edges. The separation can be solved in this setting since constrained shortest paths has a polytime algorithm if one of the two metrics is unit-cost.*

Decomposing (or clustering) graphs into lower diameter subgraphs is a scheme which occasionally arises (for instance, the tree-embedding work of [8] employs a diameter-reduction step). We close with an example in this general vein.

*Example 12 (***Diameter-Reduction).** *Suppose we would like to find a submodular-sparse cut in a graph G which decomposes it into subgraphs $G_1, G_2, \ldots$ each of which has a bounded diameter τ. We are then looking at a submodular minimization problem whose bounded blocker consists of paths of length exactly $\tau + 1$.*

4.3 Bounded Blockers and Multi-Agent Minimization

In this section we establish an $O(\beta \ln(n))$-approximation for the nonnegative monotone multi-agent submodular problem associated to families with β-bounded blockers. Notice that this result generalizes the $2\ln(n)$-approximation for multi-agent vertex cover (Goel et al [12]).

Theorem 11. *Let $\mathcal{F}$ be a family with a β-bounded blocker. Then there is a $\beta \ln(n)$-approximation algorithm for the associated multi-agent nonnegative monotone problem. If $P^*(\mathcal{F})$ has a polytime separation oracle, then this is a polytime algorithm.*

Proof. Let x^* be an optimal solution to (7). Recall that if $P^*(\mathcal{F})$ has a polytime separation oracle then we have that x^* can be computed in polynomial time (Corollary 9). Consider $Q = \{v \in V : \sum_{S \ni v} \sum_{i \in [k]} x^*(S,i) \geq \frac{1}{\beta}\}$. Since $\mathcal{F}$ has a β-bounded blocker it follows that $Q \in \mathcal{F}$, and moreover, βx^* is a fractional cover of Q. That is, $\sum_S \sum_i \beta x^*(S,i)\chi^S \geq \chi^Q$.

We now appeal to standard analysis of the greedy algorithm applied to the fractional set cover (of Q) with set costs given by $c((S,i)) = f_i(S)$. This results in an integral cover of Q whose cost is augmented by a factor of at most $\ln n$ times the fractional cover. Denote this integral cover by $(S^1,1),...,(S^{m_1},1),(S^1,2),\ldots,\ (S^{m_2},2),...,(S^{m_k},k)$. Since $\sum_S \sum_i x^*(S,i) f_i(S) \leq OPT(MA)$ we have that the integral cover cost satisfies

$$\begin{aligned}\sum_{i=1}^{k}\sum_{j=1}^{m_i} c(S^j,i) &\leq \ln(n)\sum_S\sum_i \beta x^*(S,i) f_i(S)\\ &= \beta\ln(n)\sum_S\sum_i x^*(S,i) f_i(S)\\ &\leq \beta\ln(n) OPT(MA).\end{aligned}$$

In addition, by letting $S_i = \cup_{j=1}^{m_i}(S^j,i)$, by submodularity (or just subadditivity) we have

$$\sum_{i=1}^{k} f_i(S_i) = \sum_{i=1}^{k} f_i(\cup_{j=1}^{m_i}(S^j,i)) \leq \sum_{i=1}^{k}\sum_{j=1}^{m_i} f_i((S^j,i)) = \sum_{i=1}^{k}\sum_{j=1}^{m_i} c(S^j,i).$$

Finally, if $S_i \cap S_j \neq \emptyset$ for some $i \neq j$ throw away the common shared elements from one of them arbitrarily. Denote by S_1^*,...,S_k^* these new sets, and notice that $S_i^* \subseteq S_i$ and $S_1^* + ... + S_k^* = Q$. By monotonicity of the functions f_i's we have

$$\sum_{i\in[k]} f_i(S_i^*) \leq \sum_{i\in[k]} f_i(S_i) \leq \beta\ln(n) OPT(MA)$$

and the result follows. □

5 Multi-agent minimization over $\mathcal{F} = \{V\}$

In this section we discuss the multi-agent framework (3) in the minimization setting and in the special case where the functions f_i are nonnegative and $\mathcal{F} = \{V\}$. That is, we focus on the problem

$$\begin{array}{ll}\min & \sum_{i=1}^{k} f_i(S_i)\\ \text{s.t.} & S_1 + \cdots + S_k = V\\ & S_i \in \mathcal{F}_i\,,\ \forall i \in [k]\end{array}$$

where the functions f_i are nonnegative submodular.

As previously discussed in Section 1.2 (see Example 5), the special case of this problem where the agents are allowed to take any subset of V (i.e. $\mathcal{F}_i = 2^V$ for all i) has been widely considered in the literature and it has been referred to as the Minimum Submodular Cost Allocation (MSCA) problem. We summarize the results for MSCA here.

Theorem 12. *Consider MSCA for nonnegative functions f_i. We have the following.*

1. *There is a polytime (tight) $O(\ln(n))$-approx for the case where the functions are monotone ([17,42]).*
2. *There is a polytime (tight) $O(\ln(n))$-approx for the case where the functions f_i can be written as $f_i = g_i + h$ for some nonnegative monotone submodular g_i and a nonnegative symmetric submodular function h ([3]).*
3. *There is a polytime (tight) $O(k \cdot \ln(n))$-approx for the case where the functions f_i can be written as $f_i = g_i + h$ for some nonnegative monotone submodular g_i and a nonnegative submodular function h ([7,35]).*
4. *For general functions f_i and $k \geq 3$ this problem does not admit any finite multiplicative approximation factor ([7]).*

We now discuss the monotone MSCA problem with additional constraints on the sets that the agents can take. Two of the most natural settings to consider are to impose cardinality constraints on the number of elements that an agent can take, or to allow agent i to only take elements from some subset $V_i \subseteq V$, or both. The rest of this section will be dealing with multi-agent problems that incorporate one (or both) of such constraints. We use one of these problems to provide an example of an instance in which the multivariate problem is $\Omega(n)$ hard to approximate while its "decomposable" version (i.e. MA) can be solved exactly (see Corollary 11 and Lemma 6).

The following result is an observation that derives from the work of Svitkina-Tardos [42] on the submodular facility location problem.

Corollary 10. *The problem*

$$\begin{array}{ll} \min & \sum_{i=1}^{k} f_i(S_i) \\ s.t. & S_1 + \cdots + S_k = V \\ & S_i \subseteq V_i, \end{array}$$

where the functions f_i are nonnegative monotone submodular admits a $\ln(\max_{i\in[k]} |V_i|)$-approximation.

Proof. Svitkina-Tardos [42] provide a reduction of the problem

$$\begin{array}{ll} \min & \sum_{i=1}^{k} f_i(S_i) \\ \text{s.t.} & S_1 + \cdots + S_k = V \end{array}$$

where the functions f_i are nonnegative monotone submodular to a set cover instance. Then, the greedy algorithm (using unconstrained submodular minimization as a subroutine) is used to obtain an optimal $\ln(n)$-approximation. The same reduction applied in this setting gives a t-set cover instance (i.e. an instance where all the sets have size at most t) with $t = \max_{i\in[k]} |V_i|$. Now the result follows from the fact that greedy gives a $\ln(t)$-approximation for the t-set cover problem. □

We now consider the above problem with an additional cardinality constraint on the sets the agents can take. We use the techniques from the generic reduction presented in Section 2 and a known result from matching theory in bipartite graphs to obtain the desired approximation. We first introduce some definitions that will be useful for the proof.

Let $G = (A + B, E)$ be a bipartite graph with $|A| = k$ and $A = \{a_1, ..., a_k\}$. We call a subset of edges $M \subseteq E$ *saturating* if $cov(M) = B$. Also, given some nonnegative integers $b_1, ..., b_k$, we say that M is a $(b_1, ..., b_k)$-*matching* if $|M \cap \delta(a_i)| \leq b_i$ for each $i \in [k]$.

Lemma 5. *Consider the following multi-agent minimization problem*

$$\begin{array}{ll} \min & \sum_{i\in[k]} g_i(S_i) \\ s.t. & S_1 + \cdots + S_k = V \\ & S_i \subseteq V_i \\ & |S_i| \leq b_i. \end{array}$$

Then, if the functions g_i are nonnegative monotone submodular there exists a polytime $(\max_i b_i)$-factor approximation.

Proof. Apply the generic reduction from Section 2, and let $G = ([k] + V, E)$ be the corresponding bipartite graph (i.e., with $E = \uplus_{i\in[k]}\{(i, v) : v \in V_i\}$) and f the corresponding function. Define weights $w : E \to \mathbb{R}_+$ such that $w_e = f(e)$ for each $e \in E$. Let $M \subseteq E$ be any saturating $(b_1, ..., b_k)$-matching in G. By submodularity of f we have that $w(M) \geq f(M)$. Moreover, since g can be written as $g(S_1, ..., S_k) = \sum_{i\in[k]} g_i(S_i)$, we have that $f(M) = \sum_{i\in[k]} f_i(M_i)$ for some monotone submodular functions f_i where $M_i := M \cap \delta(i)$ and $f_i(M_i) = g_i(cov(M_i))$. Moreover, by monotonicity of the f_i we have

$$f_i(M_i) \geq \max_{e\in M_i} f_i(e) = \max_{e\in M_i} f(e) \geq \frac{1}{|M_i|}\sum_{e\in M_i} f(e) = \frac{1}{|M_i|}w(M_i) \geq \frac{1}{b_i}w(M_i)$$

where the first equality follows from the fact that $f(e) = f_i(e)$ for every $e \in \delta(i)$, and the last inequality from the fact that M is a saturating $(b_1, ..., b_k)$-matching and hence $|M_i| \leq b_i$ for every $i \in [k]$. Hence, it follows that for each saturating $(b_1, ..., b_k)$-matching M we have

$$w(M) \geq f(M) = \sum_{i\in[k]} f_i(M_i) \geq \sum_{i\in[k]} \frac{1}{b_i}w(M_i) \geq \sum_{i\in[k]} (\frac{1}{\max_{i\in[k]} b_i})w(M_i) = \frac{1}{\max_{i\in[k]} b_i}w(M).$$

In particular, if M^* is a minimum saturating $(b_1, ..., b_k)$-matching for the weights w, we have

$$w(M^*) \geq f(M^*) \geq \min_{\text{s.t. } M \text{ is a saturating } (b_1,...,b_k)\text{-matching}} f(M) \geq \frac{1}{\max_{i\in[k]} b_i}w(M^*).$$

This is, $f(M^*)$ is a $(\max_i b_i)$-factor approximation. We conclude the argument by noticing that w is a modular function, and hence we can find a minimum saturating $(b_1, ..., b_k)$-matching in polynomial time. □

Corollary 11. *Consider the above problem in the special case where $k = n$ and all the $b_i = 1$. Then, an optimal solution to the problem is given by a minimum perfect matching with respect to the weights w as defined in the proof above. Moreover, this solution is still optimal even in the case where we do not require anymore the original functions g_i to be nonnegative or monotone.*

Surprisingly, the above problem becomes very hard to approximate if we allow general multivariate submodular functions instead of decomposable ones. This is true even under the special conditions of Corollary 11 in which $k = n$ and $b_i = 1$ for all the agents.

Lemma 6. *There is an information-theoretic hardness lower bound of $\Omega(n)$ for the multivariate problem*

$$\begin{array}{rc} \min & g(S_1, \ldots, S_k) \\ s.t. & S_1 + \cdots + S_k = V \\ & S_i \subseteq V_i \\ & |S_i| \leq 1 \end{array}$$

where g is a nonnegative monotone multivariate submodular function.

Proof. We prove this in the special case where $k = n$. Notice that in this case the constraint $|S_i| \leq 1$ becomes $|S_i| = 1$. We reduce an instance of the submodular perfect matching (SPM) problem to the problem above. In [12], the following is shown for any fixed $\epsilon, \delta > 0$. Any randomized $(\frac{n^{1-3\epsilon}}{1+\delta})$-approximation algorithm for the submodular minimum cost perfect matching problem on bipartite graphs (with at least one perfect matching), requires exponentially many queries. So let us consider an instance of SPM consisting of a bipartite graph $G = (A + B, E)$ with $|A| = |B|$ and a nonnegative monotone submodular function $f : 2^E \to \mathbb{R}_+$. The goal is to find a perfect matching $M \subseteq E$ of minimum submodular cost. We can transform this to the general multivariate problem above by applying the reduction from Section 2 in the opposite direction. More formally, let $k = |A|$, $V = B$, and $V_i = cov(\delta(a_i))$ where $A = \{a_1, ..., a_k\}$. Also, define a multivariate function $g : 2^{V_1} \times ... \times 2^{V_k} \to \mathbb{R}_+$ by $g(S_1, ..., S_k) = f(\uplus_{i \in [k]} \{(a_i, b) : b \in S_i\})$ for any $(S_1, ..., S_k) \in 2^{V_1} \times \cdots \times 2^{V_k}$ (i.e. $S_i \subseteq V_i$). We have g is multivariate submodular by Claim 2.2. It follows that a solution to

$$\begin{array}{rc} \min & g(S_1, \ldots, S_k) \\ \text{s.t.} & S_1 + \cdots + S_k = V \\ & S_i \subseteq V_i \\ & |S_i| = 1 \end{array}$$

(where here $|V| = k$) is a solution to the original SPM instance. □

We conclude this section by showing that monotone MSCA (which admits an $O(\ln(n))$-approximation) becomes very hard when additional cardinality constraints are added.

Claim. The problem

$$\begin{array}{rc} \min & \sum_{i=1}^{k} f_i(S_i) \\ \text{s.t.} & S_1 + \cdots + S_k = V \\ & |S_i| \leq b_i \end{array}$$

where the f_i are nonnegative monotone submodular functions is $\Omega(\frac{\sqrt{n}}{\ln n})$-hard. This is even the case when $k = 2$.

Proof. It is known ([41]) that the problem $\min f(S) : |S| = m$ where f is a nonnegative monotone submodular function is $\Omega(\frac{\sqrt{n}}{\ln n})$-hard. We reduce a general instance of this problem to the problem of the claim. To see this notice that by monotonicity of f we have

$$\begin{array}{llll} \min & f(S) & = \min \quad f(S) & = \min \qquad f(S) \qquad\quad = \min \; f(S) + f'(S') \\ \text{s.t.} & |S| = m & \text{s.t.} \; |S| \geq m & \text{s.t.} \quad S + S' = V \qquad \text{s.t.} \quad S + S' = V \\ & & & |S'| \leq |V| - m \qquad\quad |S'| \leq |V| - m \end{array}$$

where $f'(S') = 0$ for all $S' \subseteq V$. □

References

1. Buchbinder, N., Feldman, M., Seffi, J., Schwartz, R.: A tight linear time (1/2)-approximation for unconstrained submodular maximization. SIAM Journal on Computing 44(5), 1384–1402 (2015)
2. Calinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a monotone submodular function subject to a matroid constraint. SIAM Journal on Computing 40(6), 1740–1766 (2011)
3. Chekuri, C., Ene, A.: Submodular cost allocation problem and applications. In: Automata, Languages and Programming, pp. 354–366. Springer (2011)
4. Chen, C., Chester, S., Srinivasan, V., Wu, K., Thomo, A.: Group-aware weighted bipartite b-matching. In: Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. pp. 459–468. ACM (2016)
5. Chen, C., Zheng, L., Srinivasan, V., Thomo, A., Wu, K., Sukow, A.: Conflict-aware weighted bipartite b-matching and its application to e-commerce. IEEE Transactions on Knowledge and Data Engineering 28(6), 1475–1488 (2016)
6. Conforti, M., Cornuéjols, G.: Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem. Discrete applied mathematics 7(3), 251–274 (1984)
7. Ene, A., Vondrák, J.: Hardness of submodular cost allocation: Lattice matching and a simplex coloring conjecture. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014) 28, 144–159 (2014)
8. Fakcharoenphol, J., Rao, S., Talwar, K.: A tight bound on approximating arbitrary metrics by tree metrics. In: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing. pp. 448–455. ACM (2003)
9. Feige, U., Mirrokni, V.S., Vondrak, J.: Maximizing non-monotone submodular functions. SIAM Journal on Computing 40(4), 1133–1153 (2011)
10. Feldman, M., Naor, J., Schwartz, R.: A unified continuous greedy algorithm for submodular maximization. In: Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on. pp. 570–579. IEEE (2011)
11. Fisher, M.L., Nemhauser, G.L., Wolsey, L.A.: An analysis of approximations for maximizing submodular set functions-II. Springer (1978)
12. Goel, G., Karande, C., Tripathi, P., Wang, L.: Approximability of combinatorial problems with multi-agent submodular cost functions. In: Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on. pp. 755–764. IEEE (2009)
13. Goemans, M.X., Ramakrishnan, V.: Minimizing submodular functions over families of sets. Combinatorica 15(4), 499–513 (1995)
14. Goemans, M.X., Soto, J.A.: Symmetric submodular function minimization under hereditary family constraints. arXiv preprint arXiv:1007.2140 (2010)
15. Grötschel, M., Lovász, L., Schrijver, A.: Geometric algorithms and combinatorial optimization, vol. 2. Springer Science & Business Media (2012)
16. Gupta, A., Roth, A., Schoenebeck, G., Talwar, K.: Constrained non-monotone submodular maximization: Offline and secretary algorithms. In: International Workshop on Internet and Network Economics. pp. 246–257. Springer (2010)
17. Hayrapetyan, A., Swamy, C., Tardos, É.: Network design for information networks. In: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms. pp. 933–942. Society for Industrial and Applied Mathematics (2005)
18. Iwata, S., Fleischer, L., Fujishige, S.: A combinatorial strongly polynomial algorithm for minimizing submodular functions. Journal of the ACM (JACM) 48(4), 761–777 (2001)
19. Iwata, S., Nagano, K.: Submodular function minimization under covering constraints. In: Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on. pp. 671–680. IEEE (2009)
20. Iwata, S., ichi Tanigawa, S., Yoshida, Y.: Bisubmodular function maximization and extensions. Tech. rep., Technical Report METR 2013-16, The University of Tokyo (2013)

21. Khot, S., Lipton, R.J., Markakis, E., Mehta, A.: Inapproximability results for combinatorial auctions with submodular utility functions. In: International Workshop on Internet and Network Economics. pp. 92–101. Springer (2005)
22. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within 2- ε. Journal of Computer and System Sciences 74(3), 335–349 (2008)
23. Kolmogorov, V.: Submodularity on a tree: Unifying l^\ natural-convex and bisubmodular functions. In: Mathematical Foundations of Computer Science 2011, pp. 400–411. Springer (2011)
24. Krause, A., Guestrin, C.: Near-optimal observation selection using submodular functions. In: AAAI. vol. 7, pp. 1650–1654 (2007)
25. Krause, A., Leskovec, J., Guestrin, C., VanBriesen, J., Faloutsos, C.: Efficient sensor placement optimization for securing large water distribution networks. Journal of Water Resources Planning and Management 134(6), 516–526 (November 2008)
26. Krause, A., McMahan, H.B., Guestrin, C., Gupta, A.: Robust submodular observation selection. Journal of Machine Learning Research 9(Dec), 2761–2801 (2008)
27. Krause, K., Goodwin, M., Smith, R.: Optimal software test planning through automated network analysis. TRW Systems Group (1973)
28. Kulik, A., Shachnai, H., Tamir, T.: Maximizing submodular set functions subject to multiple linear constraints. In: Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 545–554. Society for Industrial and Applied Mathematics (2009)
29. Lee, J., Mirrokni, V.S., Nagarajan, V., Sviridenko, M.: Non-monotone submodular maximization under matroid and knapsack constraints. In: Proceedings of the forty-first annual ACM symposium on Theory of computing. pp. 323–332. ACM (2009)
30. Lee, J., Sviridenko, M., Vondrák, J.: Submodular maximization over multiple matroids via generalized exchange properties. Mathematics of Operations Research 35(4), 795–806 (2010)
31. Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. In: Proceedings of the 3rd ACM conference on Electronic Commerce. pp. 18–28. ACM (2001)
32. Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. Games and Economic Behavior 55(2), 270–296 (2006), http://EconPapers.repec.org/RePEc:eee:gamebe:v:55:y:2006:i:2:p:270-296
33. Lovász, L.: Submodular functions and convexity. In: Mathematical Programming The State of the Art, pp. 235–257. Springer (1983)
34. Mirrokni, V., Schapira, M., Vondrák, J.: Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In: Proceedings of the 9th ACM conference on Electronic commerce. pp. 70–77. ACM (2008)
35. Mirzakhani, M.: personal communication with Alina Ene and Jan Vondrák (2014)
36. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions - i. Mathematical Programming 14(1), 265–294 (1978)
37. Nemhauser, G.L., Wolsey, L.A.: Best algorithms for approximating the maximum of a submodular set function. Mathematics of operations research 3(3), 177–188 (1978)
38. Orlin, J.B.: A faster strongly polynomial time algorithm for submodular function minimization. Mathematical Programming 118(2), 237–251 (2009)
39. Orlin, J.B., Schulz, A.S., Udwani, R.: Robust monotone submodular function maximization. In: International Conference on Integer Programming and Combinatorial Optimization. pp. 312–324. Springer (2016)
40. Schrijver, A.: A combinatorial algorithm minimizing submodular functions in strongly polynomial time. Journal of Combinatorial Theory, Series B 80(2), 346–355 (2000)
41. Svitkina, Z., Fleischer, L.: Submodular approximation: Sampling-based algorithms and lower bounds. SIAM Journal on Computing 40(6), 1715–1737 (2011)
42. Svitkina, Z., Tardos, É.: Facility location with hierarchical facility costs. ACM Transactions on Algorithms (TALG) 6(2), 37 (2010)
43. Thapper, J., Zivny, S.: The power of linear programming for valued csps. In: Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on. pp. 669–678. IEEE (2012)

44. Vondrák, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: Proceedings of the fortieth annual ACM symposium on Theory of computing. pp. 67–74. ACM (2008)
45. Vondrák, J.: Symmetry and approximability of submodular maximization problems. SIAM Journal on Computing 42(1), 265–304 (2013)
46. Ward, J., Živnỳ, S.: Maximizing bisubmodular and k-submodular functions. In: Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 1468–1481. SIAM (2014)

APPENDIX

A Upwards-Closed (aka Blocking) Families

In this section, we give some background for blocking families. As our work for minimization is restricted to monotone functions, we can often convert an arbitrary set family into its upwards-closure (i.e., a blocking version of it) and work with it instead. We discuss this reduction as well.

A.1 Blocking Families and a Natural Relaxation for $P(\mathcal{F})$

A set family $\mathcal{F}'$, over a ground set V is *upwards-closed* if $F \subseteq F'$ and $F \in \mathcal{F}'$, implies that $F' \in \mathcal{F}'$; these are sometimes referred to as *blocking families*. Examples of such families include vertex-covers or set covers more generally, whereas spanning trees are not.

For a blocking family $\mathcal{F}'$ one normally works with the induced sub-family $\mathcal{F}$ of minimal sets. Then $\mathcal{F}$ has the property that it is a *clutter*, that is, $\mathcal{F}$ does not contain a pair of comparable sets, i.e., sets $F \subset F'$. If $\mathcal{F}$ is a clutter, then there is an associated *blocking* clutter $\mathcal{B}(\mathcal{F})$, which consists of the minimal sets B such that $B \cap F \neq \emptyset$ for each $F \in \mathcal{F}$. We refer to $\mathcal{B}(\mathcal{F})$ as the *blocker* of $\mathcal{F}$. The same definition also yields $\mathcal{B}(\mathcal{F}') = \mathcal{B}(\mathcal{F})$. The following can be checked easily.

Claim (Lehman).

1. $F \in \mathcal{F}'$ if and only if $F \cap B \neq \emptyset$ for all $B \in \mathcal{B}(\mathcal{F})$.
2. $\mathcal{B}(\mathcal{B}(\mathcal{F})) = \mathcal{F}$.

The significance of blockers is that one may assert membership in an upwards-closed family $\mathcal{F}'$ by checking intersections on sets from the blocker $\mathcal{B}(\mathcal{F})$. That is $F \in \mathcal{F}'$ if and only if $F \cap B \neq \emptyset$ for each $B \in \mathcal{B}(\mathcal{F})$. This simple observation gives a natural relaxation for exploring minimization problems over the integral polyhedron $P(\mathcal{F}') = conv(\{\chi^F : F \in \mathcal{F}'\})$. The *blocking formulation* for $\mathcal{F}'$ is:

$$P^*(\mathcal{F}') = \{z \in \mathbb{R}^V_{\geq 0} : z(B) \geq 1 \ \ \forall B \in \mathcal{B}(\mathcal{F})\}. \tag{12}$$

Clearly we have $P(\mathcal{F}') \subseteq P^*(\mathcal{F}')$.

A.2 Reducing to Blocking Families

Now consider an arbitrary set family $\mathcal{F}$ over V. We may define its *upwards closure* by $\mathcal{F}^{\uparrow} = \{F' : F \subseteq F' \text{ for some } F \in \mathcal{F}\}$. In this section we argue that in order to solve a

monotone optimization problem over sets in $\mathcal{F}$ it is often sufficient to work over its upwards-closure.

As already noted $\mathcal{B}(\mathcal{F}) = \mathcal{B}(\mathcal{F}^\uparrow)$ and hence one approach is via the blocking formulation $P^*(\mathcal{F}) = P^*(\mathcal{F}^\uparrow)$. This requires two ingredients. First, we need a separation algorithm for the blocking relaxation, but indeed this is often available for many natural families such as spanning trees, perfect matchings, *st*-paths, and vertex covers. The second ingredient needed is the ability to turn an integral solution $\chi^{F'}$ from $P^*(\mathcal{F}^\uparrow)$ or $P(\mathcal{F}^\uparrow)$ into an integral solution $\chi^F \in P(\mathcal{F})$. We now argue that this is the case if a polytime separation algorithm is available for the blocking relaxation $P^*(\mathcal{F}^\uparrow)$ or for the polytope $P(\mathcal{F}) := conv(\{\chi^F : F \in \mathcal{F}\})$.

For a polyhedron P, we denote its *dominant* by $P^\uparrow := \{z : z \geq x \text{ for some } x \in P\}$. The following observation is straightforward.

Claim. Let H be the set of vertices of the hypercube in $\mathbb{R}^V$. Then

$$H \cap P(\mathcal{F}^\uparrow) = H \cap P(\mathcal{F})^\uparrow = H \cap P^*(\mathcal{F}^\uparrow).$$

In particular we have that $\chi^S \in P(\mathcal{F})^\uparrow \iff \chi^S \in P^*(\mathcal{F}^\uparrow)$.

We can now use this observation to prove the following.

Lemma 7. *Assume we have a separation algorithm for $P^*(\mathcal{F}^\uparrow)$. Then for any $\chi^S \in P^*(\mathcal{F}^\uparrow)$ we can find in polytime $\chi^M \in P(\mathcal{F})$ such that $\chi^M \leq \chi^S$.*

Proof. Let $S = \{1, 2, \ldots, k\}$. We run the following routine until no more elements can be removed:

For $i \in S$
  If $\chi^{S-i} \in P^*(\mathcal{F}^\uparrow)$ then $S = S - i$

Let χ^M be the output. We show that $\chi^M \in P(\mathcal{F})$. Since $\chi^M \in P^*(\mathcal{F}^\uparrow)$, by Claim A.2 we know that $\chi^M \in P(\mathcal{F})^\uparrow$. Then by definition of dominant there exists $x \in P(\mathcal{F})$ such that $x \leq \chi^M \in P(\mathcal{F})^\uparrow$. It follows that the vector x can be written as $x = \sum_i \lambda_i \chi^{U_i}$ for some $U_i \in \mathcal{F}$ and $\lambda_i \in (0, 1]$ with $\sum_i \lambda_i = 1$. Clearly we must have that $U_i \subseteq M$ for all i, otherwise x would have a non-zero component outside M. In addition, if for some i we have $U_i \subsetneq M$, then there must exist some $j \in M$ such that $U_i \subseteq M - j \subsetneq M$. Hence $M - j \in \mathcal{F}^\uparrow$, and thus $\chi^{M-j} \in P(\mathcal{F})^\uparrow$ and $\chi^{M-j} \in P^*(\mathcal{F}^\uparrow)$. But then when component j was considered in the algorithm above, we would have had S such that $M \subseteq S$ and so $\chi^{S-j} \in P^*(\mathcal{F}^\uparrow)$ (that is $\chi^{S-j} \in P(\mathcal{F})^\uparrow$), and so j should have been removed from S, contradiction. □

We point out that for many natural set families $\mathcal{F}$ we can work with the relaxation $P^*(\mathcal{F}^\uparrow)$ assuming that it admits a separation algorithm. Then, if we have an algorithm which produces $\chi^{F'} \in P^*(\mathcal{F}^\uparrow)$ satisfying some approximation guarantee for a monotone problem, we can use Lemma 7 to construct in polytime $F \in \mathcal{F}$ which obeys the same guarantee.

Moreover, notice that for Lemma 7 to work we do not need an actual separation oracle for $P^*(\mathcal{F}^\uparrow)$, but rather all we need is to be able to separate over $0-1$ vectors only. Hence, since the polyhedra $P^*(\mathcal{F}^\uparrow)$, $P(\mathcal{F}^\uparrow)$ and $P(\mathcal{F})^\uparrow$ have the same $0-1$ vectors (see Claim A.2), a separation oracle for either $P(\mathcal{F}^\uparrow)$ or $P(\mathcal{F})^\uparrow$ would be enough for the routine of Lemma 7 to work. We now show that this is the case if we have a polytime separation oracle for $P(\mathcal{F})$. The following result shows that if we can separate efficiently over $P(\mathcal{F})$ then we can also separate effficiently over the dominant $P(\mathcal{F})^\uparrow$.

Claim. If we can separate over a polyhedron P in polytime, then we can also separate over its dominant $P^{\uparrow}$ in polytime.

Proof. Given a vector y, we can decide whether $y \in P^{\uparrow}$ by solving

$$\begin{aligned} x + s &= y \\ x &\in P \\ s &\geq 0. \end{aligned}$$

Since can we easily separate over the first and third constraints, and a separation oracle for P is given (i.e. we can also separate over the set of constraints imposed by the second line), it follows that we can separate over the above set of constraints in polytime. □

Now we can apply the same mechanism from Lemma 7 to turn feasible sets from $\mathcal{F}^{\uparrow}$ into feasible sets in $\mathcal{F}$.

Corollary 12. *Assume we have a separation algorithm for $P(\mathcal{F})^{\uparrow}$. Then for any $\chi^S \in P(\mathcal{F})^{\uparrow}$ we can find in polytime $\chi^M \in P(\mathcal{F})$ such that $\chi^M \leq \chi^S$.*

We conclude this section by making the remark that if we have an algorithm which produces $\chi^{F'} \in P(\mathcal{F}^{\uparrow})$ satisfying some approximation guarantee for a monotone problem, we can use Corollary 12 to construct $F \in \mathcal{F}$ which obeys the same guarantee.

B Relaxations for Constrained Submodular Minimization

Submodular optimization techniques for minimization on a set family have involved two standard relaxations, one being linear [12] and one being convex [19,3]. We introduce these in this section.

B.1 A Convex Relaxation

We will be working with upwards-closed set families $\mathcal{F}$, and their blocking relaxations $P^*(\mathcal{F})$. As we now work with arbitrary vectors $z \in [0,1]^n$, we must specify how our objective function $f(S)$ behaves on all points $z \in P^*(\mathcal{F})$. Formally, we call $g : [0,1]^V \to \mathbb{R}$ an *extension* of f if $g(\chi^S) = f(S)$ for each $S \subseteq V$. Conversely, if $g : [0,1]^V \to \mathbb{R}$ is a function, then we call its *set restriction* the function $f(S) = g(\chi^S)$ for each $S \subseteq V$.

If our relaxation is to be algorithmically useful, one anticipates the extension g to be convex. For a submodular objective function $f(S)$ there can be many convex extensions of f to $[0,1]^V$ (or to $\mathbb{R}^V$). The most popular one has been the so-called *Lovász Extension* (introduced in [33]) due to several of its desirable properties.

We give one of several equivalent definitions for the Lovász Extension. Let $0 < v_1 < v_2 < ... < v_m \leq 1$ be the distinct positive values taken in some vector $z \in [0,1]^V$. We also define $v_0 = 0$ and $v_{m+1} = 1$ (which may be equal to v_m). Define for each $i \in \{0, 1, ..., m\}$ the set $S_i = \{j : z_j > v_i\}$. In particular, S_0 is the support of z and $S_m = \emptyset$. One then defines:

$$f^L(z) = \sum_{i=0}^{m} (v_{i+1} - v_i) f(S_i).$$

Lemma 8 (Lovász [33]). *The function f^L is convex if and only if f is submodular.*

One could now attack constrained submodular minimization by solving the problem $\min f^L(z) : z \in P^*(\mathcal{F})$ and then seek rounding methods for the resulting solution. This is the approach used in [19,3].

We will later use two additional properties of the Lovász Extension. We say that a function g is (positively) *homogeneous* if for each $c \geq 0$, we have $g(cz) = cg(z)$.

Lemma 9. *The Lovász Extension of f is homogeneous if f is submodular and normalized (i.e. $f(\emptyset) = 0$). That is for any $x \in [0,1]^n$ and $\alpha \in [0,1]$ $f^L(\alpha x) = \alpha f^L(x)$. In particular, for any set Q we have $f^L((\alpha)\chi^Q) = \alpha f(Q)$.*

Proof. This follows by construction. Recall that

$$f^L(x) = \sum_{i=0}^{m} (v_{i+1} - v_i) f(S_i)$$

where $v_0 = 0$, $v_{m+1} = 1$, and $0 \leq v_1 < v_2 < ... < v_m \leq 1$ are the distinct values taken in the vector $x \in [0,1]^n$. The sets S_i are defined as $S_i = \{j : x_j > v_i\}$ for each $i \in \{0, 1, ..., m\}$. It follows directly that

$$f^L(\alpha x) = \sum_{i=0}^{m-1} (\alpha)(v_{i+1} - v_i) f(S_i) + (1 - \alpha v_m) f(\emptyset) = \alpha f^L(x)$$

□

Lemma 10. *If $f(S)$ is a monotonic set function, then f^L is a monotonic function.*

This can be used to establish the following.

Corollary 13. *Let $g : [0,1]^V \to \mathbb{R}$ be a convex, monotonic, homogeneous function. Suppose that $z \leq \sum_S x(S)\chi^S$, then $g(z) \leq \sum_S x(S) g(\chi^S)$. In particular, if f^L is the extension of a normalized, monotonic, submodular function $f(S)$ and $\chi^Z \leq \sum_S x(S)\chi^S$, then $f(Z) \leq \sum_S x(S) f(S)$.*

Proof. We prove the latter statement, the first being similar. Without loss of generality, $C := \sum_{S \subseteq V} x(S) > 0$. Hence:

$$\begin{aligned}
f(Q) &= f^L(\chi^Q) && \text{(by definition)} \\
&= \tfrac{C}{C} f^L(\chi^Q) \\
&= C f^L(\tfrac{1}{C}\chi^Q) && \text{(by homogeneity)} \\
&\leq C f^L(\tfrac{1}{C}\textstyle\sum_{S \subseteq V} x(S)\chi^S) && \text{(by monotonicity)} \\
&\leq C \textstyle\sum_{S \subseteq V} (\frac{x(S)}{C}) f^L(\chi^S) && \text{(by convexity)} \\
&= \textstyle\sum_{S \subseteq V} x(S) f^L(\chi^S) \\
&= \textstyle\sum_{S \subseteq V} x(S) f(S). && \text{(by definition).}
\end{aligned}$$

□

B.2 An Extended LP Formulation

By Claim A.1 (see Appendix A), an integer programming formulation for the single-agent problem is

$$\begin{aligned}
\min \quad & \textstyle\sum_{S \subseteq V} f(S) x(S) \\
\text{s.t.} \quad & \textstyle\sum_{v \in B} \sum_{S \ni v} x(S) \geq 1, \ \forall B \in \mathcal{B}(\mathcal{F}) \\
& x(S) \geq 0.
\end{aligned} \tag{13}$$

The corresponding dual is given by

$$\begin{array}{ll} \max & \sum_{B\in\mathcal{B}(\mathcal{F})} y_B \\ \text{s.t.} & \sum_{v\in S}\sum_{B\ni v} y_B \leq f(S),\ \forall S\subseteq V \\ & y_B \geq 0. \end{array} \tag{14}$$

Using the ideas from [12], one may solve these LPs in polytime if the blocking family is not too large.

Lemma 11. *If f is submodular and $|\mathcal{B}(\mathcal{F})| \in poly(n)$, then the linear programs (13,14) can be solved in polytime.*

Proof. It is well-known [15] that a class of LPs can be solved in polytime if (i) it admits a polytime separation algorithm and (ii) has a polynomially bounded number of variables. Moreover, if this holds, then the class of dual problems can also be solved in polytime. Hence we focus on the dual (14) for which condition (ii) is assured by our assumption on $|\mathcal{B}(\mathcal{F})|$. Let's now consider its separation problem. For a fixed vector $y \geq 0$ we define $z_y(S) := \sum_{v\in S}\sum_{v\in B} y_B$ for any $S \subseteq V$. Notice that z_y is a modular function. Hence, $f - z_y$ is submodular and so we can solve the problem $\min_{S\subseteq V} f(S) - z_y(S)$ exactly in polytime. It follows that y is feasible if and only if this minimum is nonnegative. Thus, we can separate in the dual efficiently. □

The restriction that $|\mathcal{B}(\mathcal{F})| \in poly(n)$ is quite severe and we now show how it may be weakened.

B.3 A Combined LP and Convex Program

In this section we show that one may solve the extended LP approximately as long as a polytime separation algorithm for $P^*(\mathcal{F})$ is available. To do this, we work with a combination of the LP and the convex relaxation from the prior section. For a feasible solution x to (13), we define its *image* as

$$z = \sum_S x(S)\chi^S.$$

We sometimes work with an extended formulation

$$Q = \{(x, z) : z \text{ is the image of feasible } x \text{ for (13)}\}.$$

One establishes easily that $P^*(\mathcal{F})$ is just a projection of Q.

Lemma 12. $P^*(\mathcal{F}) = \{z : \exists x \text{ such that } (x, z) \in Q\}$.

In fact, the convex and LP relaxations are equivalent in the following sense, due to the pleasant properties of the Lovász Extension.

Lemma 13.

$$\min\{f^L(z) : z \in P^*(\mathcal{F})\} = \min\{f^L(z) : (x, z) \in Q\} = \min\{\sum_S x(S)f(S) : (x, z) \in Q\}.$$

Moreover, given a feasible solution $z^* \in P^*(\mathcal{F})$ we may compute x^* such that $(x^*, z^*) \in Q$ and $\sum_S x^*(S)f(S) = f^L(z^*)$. This may be done in time polynomially bounded in n and the encoding size of z^*.

Proof. The first equality follows from the preceding lemma. For the second equality we consider $(x, z) \in Q$. By Corollary 13, we have that $f^L(z) \leq \sum_S x(S)f(S)$. Hence $\leq$ holds. Conversely, suppose that z^* is an optimal solution for the convex optimization problem and let x^* be determined by the "level sets" associated with z^* in the definition of f^L. By definition we have $\sum_S x^*(S)f(S) = f^L(z^*)$. Moreover, one easily checks that since $z^*(B) \geq 1$ for each blocking set B, we also have that x^* is feasible for (13). Hence $\geq$ holds as well. □

Polytime Algorithms. One may apply the Ellipsoid Method to obtain a polytime algorithm which approximately minimizes a convex function over a polyhedron K as long as various technical conditions hold. For instance, one could require that there are two ellipsoids $E(a, A) \subseteq K \subseteq E(a, B)$ whose encoding descriptions are polynomially bounded in the input size for K. We should also have polytime (or oracle) access to the convex objective function defined over $\mathbf{R}^n$. In addition, one must be able to polytime solve the subgradient problem for f.[5] One may check that the subgradient problem is efficiently solvable for Lovász extensions of polynomially encodable submodular functions. We call f *polynomially encodable* if the values $f(S)$ have encoding size bounded by a polynomial in n(we always assume this for our functions). If these conditions hold, then methods from [15] imply that for any $\epsilon > 0$ we may find an approximately feasible solution for K which is approximately optimal. By approximate here we mean for instance that the objective value is within ϵ of the real optimum. This can be done in time polynomially bounded in n (size of input say) and $\log \frac{1}{\epsilon}$. Let us give a few details for our application.

In our setting, the rightmost LP from the previous lemma is essentially identical to the linear program (13); it just works over Q which is like carrying some extra meaningless variables z. The preceding result says that we can work instead over the more compact (in terms of variables) space $P^*(\mathcal{F})$ at the price of using a convex objective function. Our convex problem's feasible space is $P^*(\mathcal{F})$ and it is easy to verify that our optimal solutions will lie in the $0-1$ hypercube H. So we may define the feasible space to be H and the objective function to be $g(x) = f^L(x)$ if $x \in H \cap P^*(\mathcal{F})$ and $= \infty$ otherwise. (Clearly g is convex in $\mathbf{R}^n$ since it is a pointwise maximum of two convex functions; alternatively, one may define the Lovász Extension on $\mathbf{R}^n$ which is also fine.) Note that g can be evaluated in polytime by the definition of f^L as long as f is polynomially encodable. We can now easily find an ellipsoid inside H and one containing H each of which has poly encoding size. We may thus solve the convex problem to within $\pm\epsilon$-optimality in time bounded by a polynomial in n and $\log \frac{1}{\epsilon}$. Note that our lemmas guarantee to produce $(x^*, z^*) \in Q$ so we have exact feasibility.

One final comment about exact solvability of these problems. Since our convex problem is equivalent to an LP (over Q with nonzero coefficients on the x-variables), we know that an optimal solution (x^*, z^*) will be determined at a vertex of the polyhedron Q. It is straightforward to check by contradiction that z^* is in fact a vertex of $P^*(\mathcal{F})$. As any such vertex is defined by an $n \times n$ nonsingular system, and hence by Cramer's Rule, the encoding size of z^* is polynomially bounded (in n). But then x^* could also be viewed as a vertex to a polyhedron of the form: $\{x : Ax = z^*, x \geq 0\}$. The support of any such vertex is again determined by a nonsingular linear system with at most n rows. As z^* is polybounded, so is the encoding size of x^*. Since f and x^* are *polynomially-encodable*, we have the the optimal solution to the rightmost LP (and hence all problems) is polynomially bounded. Hence we can pre-select a polybounded $\epsilon > 0$ which guarantees us to find the exact optimum, not just an ϵ-optimal one. Thus we have the following.

[5] For a given y, find a subgradient of f at y.

Corollary 14. *Consider a class of problems $\mathcal{F}, f$ for which f's are submodular and polynomially-encodable in $n = |V|$. If there is a polytime separation algorithm for the family of polyhedra $P^*(\mathcal{F})$, then one may find an exact optimum z^* for the convex program in Lemma 13 in polytime. By Lemma 13 we may also produce x^* in polytime for which $(x^*, z^*) \in Q$ and $\sum_S x^*(S) f(S) = f^L(z^*)$.*

B.4 The Multi-Agent Formulations

The formulations above have natural extensions to the multi-agent setting. In the following, we only work with the extension of the LPs (13,14). This was already introduced in [12].

$$\begin{array}{ll} \min & \sum_{S \subseteq V} \sum_{i \in [k]} f_i(S) x(S,i) \\ \text{s.t.} & \sum_{v \in B} \sum_{S \ni v} \sum_{i \in [k]} x(S,i) \geq 1, \, \forall B \in \mathcal{B}(\mathcal{F}) \\ & x(S,i) \geq 0. \end{array} \tag{15}$$

The dual is:

$$\begin{array}{ll} \max & \sum_{B \in \mathcal{B}(\mathcal{F})} y_B \\ \text{s.t.} & \sum_{v \in S} \sum_{B \ni v} y_B \leq f_i(S), \, \forall i, S \subseteq V \\ & y_B \geq 0. \end{array} \tag{16}$$

Again, using the ideas from Goel et al. [12], one sees that separation over the dual amounts to k submodular minimization problems (cf. Lemma 11). Hence if $\mathcal{B}(\mathcal{F})$ is poly-bounded, we may solve the LP exactly.

We can also solve it in polytime, however, if we have polytime separation of $P^*(\mathcal{F})$. This follows the approach from the previous section (see Lemma 13 and Corollary 14) except our convex program now has k vectors of variables $z_1, z_2, \ldots, z_k$ (one for each agent) such that $z = \sum_i z_i$. We force the constraint $z \in P^*(\mathcal{F})$. We also have constraints forcing that each z_i is the image of an associated $(x_i(S))$ vector, i.e. $z_i = \sum_S x(S,i) \chi^S$. Then it is straightforward to check that for any feasible $z^* = z_1^* + \cdots + z_k^* \in P^*(\mathcal{F})$ we can find in polytime an x^* that is feasible for (15) and such that $\sum_i f_i^L(z_i^*) = \sum_i \sum_S f_i(S) x^*(S,i)$. These lead to the following two corollaries.

Corollary 15. *Assume there is a polytime separation oracle for $P^*(\mathcal{F})$. Then we can solve the multi-agent LP described by (15) in polytime.*

Corollary 16. *Assume we can solve the single-agent LP described by (13) in polytime. Then we can also solve the multi-agent LP described by (15) in polytime.*

Proof. If we can solve (13) then we can also solve the convex program from Lemma 13. Hence, we can separate over $P^*(\mathcal{F})$ in polynomial time. Now the statement follows from Corollary 15. □

C Properties of multivariate submodular functions

In this section we discuss several properties of multivariate submodular functions. We see that some of the characterizations and results that hold for univariate submodular functions extend naturally to the multivariate setting.

We start by showing that our definition of submodularity in the multivariate setting captures the diminishing return property. Recall that we usually think of the pair $(i, v) \in [k] \times V$ as the assignment of element v to agent i. We use this to introduce some notation for adding an element to a tuple.

Definition 1. *Given a tuple* $(S_1, \ldots, S_k) \in (2^V)^k$ *and* $(i, v) \in [k] \times V$*, we denote by* $(S_1, \ldots, S_k) + (i, v)$ *the new tuple* $(S_1, \ldots, S_{i-1}, S_i + v, S_{i+1}, \ldots, S_k)$.

Then, it is natural to think of the quantity

$$f(S_1, \ldots, S_{i-1}, S_i + v, S_{i+1}, \ldots, S_k) - f(S_1, \ldots, S_{i-1}, S_i, S_{i+1}, \ldots, S_k) \tag{17}$$

as the marginal gain of assigning element v to agent i in the tuple $(S_1, \ldots, S_k)$. Notice that with the notation introduced in Definition 1 we have that (17) can be also written as

$$f((S_1, \ldots, S_k) + (i, v)) - f(S_1, \ldots, S_k).$$

This leads to the following diminishing returns characterizations in the multivariate setting.

Proposition 3. *A multivariate function* $f : (2^V)^k \to \mathbb{R}$ *is multi-submodular if and only if for all tuples* $(S_1, \ldots, S_k) \subseteq (T_1, \ldots, T_k)$ *and* $(i, v) \in [k] \times V$ *we have*

$$f((S_1, \ldots, S_k) + (i, v)) - f(S_1, \ldots, S_k) \geq f((T_1, \ldots, T_k) + (i, v)) - f(T_1, \ldots, T_k). \tag{18}$$

Proof. We make use of the lifting reduction presented in Section 2.1. Let $\pi : (2^V)^k \to 2^E$ denote the bijection discussed in Section 2.1, and let $f' : 2^E \to \mathbb{R}$ be the univariate function defined by $f'(S) = f(\pi^{-1}(S))$. Also, let $S := \pi(S_1, \ldots, S_k)$ and $T := \pi(T_1, \ldots, T_k)$. Since $(S_1, \ldots, S_k) \subseteq (T_1, \ldots, T_k)$ we know that $S \subseteq T$. Moreover, notice that

$$f((S_1, \ldots, S_k) + (i, v)) - f(S_1, \ldots, S_k) = f'(S + (i, v)) - f'(S)$$

and

$$f((T_1, \ldots, T_k) + (i, v)) - f(T_1, \ldots, T_k) = f'(T + (i, v)) - f'(T).$$

In addition, from Claim 2.2 in Section 2.2 we know that f is multi-submodular if and only if f' is submodular. Then the result follows by observing the following.

$$\begin{aligned}
& f \text{ multi-submodular} \\
\iff & f' \text{ submodular} \\
\iff & f'(S + (i, v)) - f'(S) \geq f'(T + (i, v)) - f'(T) \quad \text{for all } S \subseteq T \text{ and } (i, v) \in E \\
\iff & f((S_1, \ldots, S_k) + (i, v)) - f(S_1, \ldots, S_k) \geq f((T_1, \ldots, T_k) + (i, v)) - f(T_1, \ldots, T_k) \\
& \quad \text{for all } (S_1, \ldots, S_k) \subseteq (T_1, \ldots, T_k) \text{ and } (i, v) \in [k] \times V.
\end{aligned}$$

□

The proof from the above proposition also shows the following characterization for multivariate submodular functions.

Proposition 4. *A multivariate function* $f : (2^V)^k \to \mathbb{R}$ *is multi-submodular if and only if for all tuples* $(S_1, \ldots, S_k)$ *and* $(i, v), (j, u) \in [k] \times V$ *we have*

$$f((S_1, \ldots, S_k) + (i, v)) - f(S_1, \ldots, S_k) \geq f((S_1, \ldots, S_k) + (j, u) + (i, v)) - f((S_1, \ldots, S_k) + (j, u)). \tag{19}$$

We now provide an explicit example of a multi-submodular function that leads to interesting applications.

Lemma 14. *Consider a quadratic function $h : \mathbb{Z}_+^k \to \mathbb{R}$ given by $h(z) = z^T Az$ for some matrix $A = (a_{ij})$. Let $f : (2^V)^k \to \mathbb{R}$ be a multivariate set function defined as $f(S_1, \dots, S_k) = h(|S_1|, \dots, |S_k|)$. Then, f is multi-submodular if and only if A satisfies*

$$a_{ij} + a_{ji} \leq 0 \quad \forall i, j \in [k]. \tag{20}$$

Proof. By Proposition 4 we know that f is multi-submodular if and only if condition (19) is satisfied. Let $(S_1, \dots, S_k)$ be an arbitrary tuple and let $(i, v), (j, u) \in [k] \times V$ such that $v \notin S_i, u \notin S_j$. Denote by z^0 the integer vector with components $z_i^0 = |S_i|$. That is, $z^0 = (|S_1|, |S_2|, \dots, |S_k|) \in \mathbb{Z}_+^k$. We call z^0 the cardinality vector associated to the tuple $(S_1, \dots, S_k)$. In a similar way, let z^1 be the cardinality vector associated to the tuple $(S_1, \dots, S_k) + (i, v)$, z^2 the cardinality vector associated to $(S_1, \dots, S_k) + (j, u)$, and z^3 the cardinality vector associated to $(S_1, \dots, S_k) + (i, v) + (j, u)$. Now notice that condition (19) can be written as

$$h(z^1) - h(z^0) \geq h(z^3) - h(z^2) \tag{21}$$

for all $z^0, z^1, z^2, z^3 \in \mathbb{Z}_+^k$ such that $z^1 = z^0 + \mathbf{e_i}$, $z^2 = z^0 + \mathbf{e_j}$, and $z^3 = z^0 + \mathbf{e_i} + \mathbf{e_j}$, where $\mathbf{e_i}$ is the vector with all components equal to zero except the ith component which equals 1, and similarly for $\mathbf{e_j}$.

We show that (21) is equivalent to (20). First notice that for a vector $z = (z_1, \dots, z_k)$ the function h can be written as $h(z) = \sum_{l,m=1}^k a_{lm} z_l z_m$. Then, using that $z_l^1 = z_l^0$ for all $l \neq i$ and $z_i^1 = z_1^0 + 1$, we have

$$h(z^1) - h(z^0) = \sum_{l,m=1}^k a_{lm} z_l^1 z_m^1 - \sum_{l,m=1}^k a_{lm} z_l^0 z_m^0 = \sum_{l=1}^k a_{li} z_l^0 + \sum_{m=1}^k a_{im} z_m^0 + a_{ii}.$$

Similarly, using that $z_l^3 = z_l^2$ for all $l \neq i$ and $z_i^3 = z_i^2 + 1$, we have

$$h(z^3) - h(z^2) = \sum_{l=1}^k a_{li} z_l^2 + \sum_{m=1}^k a_{im} z_m^2 + a_{ii}.$$

Thus, using that $z^2 = z^0 + \mathbf{e_j}$ we get

$$\begin{aligned}
& h(z^1) - h(z^0) \geq h(z^3) - h(z^2) \\
\iff & \sum_{l=1}^k a_{li} z_l^0 + \sum_{m=1}^k a_{im} z_m^0 + a_{ii} \;\geq\; \sum_{l=1}^k a_{li} z_l^2 + \sum_{m=1}^k a_{im} z_m^2 + a_{ii} \\
\iff & \sum_{l=1}^k a_{li} (z_l^0 - z_l^2) + \sum_{m=1}^k a_{im} (z_m^0 - z_m^2) \geq 0 \\
\iff & -a_{ji} - a_{ij} \geq 0 \\
\iff & a_{ji} + a_{ij} \leq 0.
\end{aligned}$$

□